

EXPRESS MAIL MAILING LABEL NUMBER EV 318623684 US

**PATENT**

**10830.0096.NPUS00**

**EMC-03-028**

**APPLICATION FOR UNITED STATES LETTERS PATENT**

**for**

**DATA RECOVERY WITH INTERNET PROTOCOL REPLICATION  
WITH OR WITHOUT FULL RESYNC**

**By**

**Philippe Armangau  
Milena Bergant  
Karthikeyan Srinivasan**

## BACKGROUND OF THE INVENTION

### 1. Limited Copyright Waiver

A portion of the disclosure of this patent document contains computer code listings and command formats to which the claim of copyright protection is made. The copyright owner has no objection to the facsimile reproduction by any person of the patent document or the patent disclosure, as it appears in the U.S. Patent and Trademark Office patent file or records, but reserves all other rights whatsoever.

### 2. Field of the Invention

The present invention relates generally to data storage systems, and more particularly to data recovery in a data storage system using remote data replication.

### 3. Description of the Related Art

Remote copy systems have been used for automatically providing data backup at a remote site in order to insure continued data availability after a disaster at a primary site. Such a remote copy facility is described in Ofek, U.S. Patent 5,901,327 issued May 4, 1999, entitled "Bundling of Write Data from Channel Commands in a Command Chain for Transmission over a Data Link Between Data Storage Systems For Remote Data Mirroring," incorporated herein by reference. This remote copy facility uses a dedicated network link and a link-layer protocol for 1:1 replication between a primary storage system and a secondary storage system. This kind of remote copy system is relatively expensive, however, because a secondary storage system and a network link are dedicated for backup purposes.

## SUMMARY OF THE INVENTION

In accordance with one aspect of the invention, there is provided a method of recovery in an asynchronous remote copy system having a primary file system at a primary site and a secondary file system at a secondary site. The primary site becomes inoperative during read/write access to the primary file system and asynchronous remote copy of changes made to the primary file system being copied to the secondary file system. The method includes responding to the primary site becoming inoperative by beginning read/write access to the secondary file system, making a snapshot copy of the secondary file system at the beginning of read/write access to the secondary file system, and keeping a record of changes made to the secondary file system during the read/write access to the secondary file system. Thereafter, when the primary site becomes operative, the snapshot copy is used to restore the primary file system to the state of the secondary file system existing when read/write access of the secondary file system was begun, and then the changes made to the secondary file system during the read/write access to the secondary file system are written into the primary file system. The method further includes terminating read/write access to the secondary file system, and once the changes made to the secondary file system have been written into the primary file system, restarting the read/write access to the primary file system and asynchronous remote copy of changes made to the primary file system being copied to the secondary file system.

In accordance with another aspect, the invention provides a method of recovery in an asynchronous remote copy system having a primary file system at a primary site and a secondary file system at a secondary site. The primary site becomes inoperative during read/write access to the primary file system and asynchronous remote copy of changes

1 made to the primary file system being copied to the secondary file system by beginning  
2 read/write access to the secondary file system. The method includes responding to the  
3 primary site becoming operative by synchronizing the primary file system to the  
4 secondary file system by beginning asynchronous remote copy of changes made to the  
5 secondary file system being copied to the primary file system, making a snapshot copy of  
6 the secondary file system at the beginning of the asynchronous remote copy of changes  
7 made to the secondary file system being copied to the primary file system, synchronizing  
8 the primary file system to the snapshot copy of the secondary file system, and once the  
9 primary file system has been synchronized to the snapshot copy of the secondary file  
10 system, terminating read/write access to the secondary file system and completing the  
11 asynchronous remote copy of changes made to the secondary file system being copied to  
12 the primary file system. The method further includes restarting the read/write access to  
13 the primary file system and asynchronous remote copy of changes made to the primary  
14 file system being copied to the secondary file system once the asynchronous remote copy  
15 of changes made to the secondary file system being copied into the primary file system  
16 has been completed.

17 In accordance with another aspect of the invention, there is provided a method of  
18 recovery in an asynchronous remote copy system having a primary file system at a  
19 primary site and a secondary file system at a secondary site. The method includes  
20 responding to the primary site becoming inoperative during read/write access to the  
21 primary file system and asynchronous remote copy of changes made to the primary file  
22 system being copied to the secondary file system by beginning read/write access to the  
23 secondary file system, making a first snapshot copy of the secondary file system at the

1 beginning of read/write access to the secondary file system, and keeping a record of  
2 changes made to the secondary file system during the read/write access to the secondary  
3 file system. When the primary site becomes operative, the method further includes  
4 making a second snapshot copy of the secondary file system, beginning asynchronous  
5 remote copy of changes made to the secondary file system since the second snapshot  
6 copy being copied to the primary file system, using the first snapshot copy to restore the  
7 primary file system to the state of the secondary file system existing when read/write  
8 access of the secondary file system was begun, and then writing into the primary file  
9 system the changes made to the secondary file system during the read/write access to the  
10 secondary file system between the time of the first snapshot copy and the second  
11 snapshot copy. The method further includes terminating read/write access to the  
12 secondary file system, and once the changes made to the secondary file system have been  
13 written into the primary file system, restarting read/write access to the primary file  
14 system and asynchronous remote copy of changes made to the primary file system being  
15 copied to the secondary file system.

16 In accordance with another aspect of the invention, there is provided a method of  
17 recovery from a disruption at a primary site in an asynchronous remote copy system in  
18 which changes made to data blocks of a primary file system at the primary site are  
19 transmitted to a secondary file system at a secondary site. The primary site stores a list of  
20 the data blocks that have been changed in the primary file system. The method includes  
21 accessing the list of the data blocks that have been changed in the primary file system to  
22 restore the primary file system to a prior state at a restart point. The prior state at the  
23 restart point includes changes made to the primary file system that have been transmitted

1 to the secondary site. The primary file system is restored by determining from the list the  
2 data blocks that have been changed in the primary file system since the restart point, and  
3 obtaining from the secondary site the data existing at the time of the restart point in the  
4 data blocks that have been changed in the primary file system since the restart point, and  
5 writing into the primary file system the data existing at the time of the restart point in the  
6 data blocks that have been changed in the primary file system since the restart point.

7 In accordance with yet another aspect, the invention provides an asynchronous  
8 remote copy system including a primary data storage system and a secondary data storage  
9 system. The primary data storage system has a primary file system, and the secondary  
10 data storage system has a secondary file system. The primary data storage system is  
11 programmed for read/write access to the primary file system and asynchronous remote  
12 copy of changes made to the primary file system being copied to the secondary file  
13 system. The secondary data storage system is programmed to respond to the primary  
14 data storage system becoming inoperative during the asynchronous remote copy of  
15 changes made to the primary file system being copied to the secondary file system by  
16 beginning read/write access to the secondary file system, making a snapshot copy of the  
17 secondary file system at the beginning of read/write access to the secondary file system,  
18 and keeping a record of changes made to the secondary file system during the read/write  
19 access to the secondary file system. Moreover, the primary data storage system and the  
20 secondary data storage system are programmed for recovery when the primary data  
21 storage system becomes operative by using the snapshot copy to restore the primary file  
22 system to the state of the secondary file system existing when read/write access of the  
23 secondary file system was begun, and then writing into the primary file system the

1 changes made to the secondary file system during the read/write access to the secondary  
2 file system, terminating read/write access to the secondary file system, and once the  
3 changes made to the secondary file system have been written into the primary file system,  
4 restarting read/write access to the primary file system and asynchronous remote copy of  
5 changes made to the primary file system being copied to the secondary file system.

6 In accordance with another aspect, the invention provides an asynchronous  
7 remote copy system including a primary data storage system and a secondary data storage  
8 system. The primary data storage system has a primary file system and the secondary  
9 data storage system has a secondary file system. The primary data storage system is  
10 programmed for read/write access to the primary file system and asynchronous remote  
11 copy of changes made to the primary file system being copied to the secondary file  
12 system. The secondary data storage system is programmed to respond to the primary  
13 data storage system becoming inoperative during the asynchronous remote copy of  
14 changes made to the primary file system being copied to the secondary file system by the  
15 secondary data storage system beginning read/write access to the secondary file system.  
16 The primary data storage system and the secondary data storage system are programmed  
17 to respond to the primary data storage system becoming operative by synchronizing the  
18 primary file system to the secondary file system by beginning asynchronous remote copy  
19 of changes made to the secondary file system being copied to the primary file system,  
20 making a snapshot copy of the secondary file system at the beginning of the  
21 asynchronous remote copy of changes made to the secondary file system being copied to  
22 the primary file system, synchronizing the primary file system to the snapshot copy of the  
23 secondary file system, and once the primary file system has been synchronized to the

1 snapshot copy of the secondary file system, terminating read/write access to the  
2 secondary file system and completing the asynchronous remote copy of changes made to  
3 the secondary file system being copied to the primary file system, and once the  
4 asynchronous remote copy of changes made to the secondary file system has been  
5 completed, restarting read/write access to the primary file system and asynchronous  
6 remote copy of changes made to the primary file system being copied to the secondary  
7 file system.

8 In accordance with yet another aspect, the invention provides an asynchronous  
9 remote copy system including a primary data storage system and a secondary data storage  
10 system. The primary data storage system has a primary file system, and the secondary  
11 data storage system has a secondary file system. The primary data storage system is  
12 programmed for read/write access to the primary file system and asynchronous remote  
13 copy of changes made to the primary file system being copied to the secondary file  
14 system. The secondary data storage system is programmed to respond to the primary  
15 data storage system becoming inoperative during the asynchronous remote copy of  
16 changes made to the primary file system being copied to the secondary file system by  
17 beginning read/write access to the secondary file system, making a first snapshot copy of  
18 the secondary file system at the beginning of read/write access to the secondary file  
19 system, and keeping a record of changes made to the secondary file system during the  
20 read/write access to the secondary file system. The secondary data storage system is  
21 further programmed to respond to the primary data storage system becoming operative by  
22 making a second snapshot copy of the secondary file system, and beginning  
23 asynchronous remote copy of changes made to the secondary file system since the second

1 snapshot copy being copied to the primary file system. Moreover, the primary data  
2 storage system and the secondary data storage system are programmed for using the first  
3 snapshot copy to restore the primary file system to the state of the secondary file system  
4 existing when read/write access of the secondary file system was begun, and then writing  
5 into the primary file system the changes made to the secondary file system during the  
6 read/write access to the secondary file system between the time of the first snapshot copy  
7 and the second snapshot copy, and terminating read/write access to the secondary file  
8 system, and once the changes made to the secondary file system have been written into  
9 the primary file system, restarting read/write access to the primary file system and  
10 asynchronous remote copy of changes made to the primary file system being copied to  
11 the secondary file system.

## 12 13 **BRIEF DESCRIPTION OF THE DRAWINGS**

14 Other objects and advantages of the invention will become apparent upon reading  
15 the following detailed description with reference to the accompanying drawings wherein:

16 FIG. 1 is a block diagram of an IP network including multiple hosts and multiple  
17 data mover computers;

18 FIG. 2 is a block diagram showing a primary data mover distributing remote copy  
19 data to multiple secondary data movers in the IP network by establishing a Transmission  
20 Control Protocol (TCP) connection with each of the secondary data movers;

21 FIG. 3 is a block diagram showing a primary data mover distributing remote copy  
22 data to multiple data movers through forwarder data movers;

1           FIG. 4 is a block diagram showing a shared save volume used to buffer local copy  
2 data transmitted from a primary data mover to a secondary data mover;

3           FIG. 5 is a block diagram showing a primary save volume and a secondary save  
4 volume;

5           FIG. 6 is a flowchart showing local replication in the system of FIG. 4;

6           FIG. 7 is a flowchart showing remote replication in the system of FIG. 5;

7           FIG. 8 is a block diagram of a primary site, including layered programming in a  
8 primary data mover;

9           FIG. 9 is a block diagram of a secondary site, including layered programming in a  
10 secondary data mover;

11          FIG. 10 is a flowchart of a process of replication at the primary site of FIG. 8;

12          FIG. 11 is a flowchart of a procedure for producing a new remote copy of a  
13 primary file system concurrent with ongoing replication and multicasting of  
14 modifications to the primary file system;

15          FIG. 12 is a flowchart of an IP-replication send-thread introduced in FIG. 8;

16          FIG. 13 is a block diagram of a volume multicast level in the data mover  
17 programming of FIG. 8 and FIG. 9;

18          FIG. 14 is a block diagram of the RCP level in the primary data mover  
19 programming of FIG. 8;

20          FIG. 15 is a block diagram of the RCP level in the secondary data mover  
21 programming of FIG. 9;

22          FIG. 16 is a block diagram of a RPC forwarder at the RPC level in a forwarder  
23 data mover;

1 FIG. 17 is a flowchart of an inbound RCP session in the secondary data mover;

2 FIG. 18 is a block diagram showing a forwarder data mover performing local  
3 replication;

4 FIG. 19 is a block diagram showing the sharing of a data mover's single TCP port  
5 for RCP connections with Hypertext Transfer Protocol (HTTP) connections;

6 FIG. 20 is a block diagram showing the state of a remote replication system  
7 before a disaster at the primary site;

8 FIG. 21 is a flowchart of a failover and recovery method performed upon the  
9 remote replication system of FIG. 20 in order to recover from the disaster at the primary  
10 site;

11 FIG. 22 is a block diagram showing the state of the remote replication system of  
12 FIG. 20 after failover from the primary site to the secondary site;

13 FIG. 23 is a flowchart of a subroutine of failover with a checkpoint and without  
14 sync, as used in the flowchart of FIG. 21;

15 FIG. 24 shows a block map that can be used to create a snapshot copy of the  
16 secondary file system at a restart point during the failover of FIG. 23;

17 FIG. 25 shows a snapshot copy facility that can be used to create the snapshot  
18 copy of the secondary file system at the restart point during the failover of FIG. 23;

19 FIG. 26 is a flowchart of a procedure for writing a block of data to a production  
20 file system in the snapshot copy facility of FIG. 25;

21 FIG. 27 is a flowchart of a procedure for reading a block of data from a snapshot  
22 file system in the snapshot copy facility of FIG. 25;

1           FIG. 28 is a flowchart of a subroutine to resync the primary file system with the  
2 secondary file system, as used in the flowchart of FIG. 21;

3           FIG. 29 shows the state of the remote replication system of FIG. 20 during the  
4 resync procedure;

5           FIG. 30 is a flowchart of a subroutine to failback to the primary file system, as  
6 used in the flowchart of FIG. 21;

7           FIG. 31 is a flowchart of execution of a failover command having a sync option  
8 and a checkpoint option;

9           FIG. 32 is a flowchart of a subroutine for failover without sync, as used in the  
10 flowchart of FIG. 31;

11          FIG. 33 is a subroutine for failover with sync, as used in the flowchart of FIG. 31;

12          FIGS. 34 and 35 comprise a flowchart of a procedure for copying snapshot delta  
13 for snapshots L to M;

14          FIGS. 36 shows a block diagram of a data network in which snapshot deltas are  
15 transmitted over a wide-area network from a network file server to a local file server in  
16 order to update the local file system as needed;

17          FIGS. 37 and 38 comprise a flowchart of a procedure for replicating snapshots in  
18 the data network of FIG. 36;

19          FIG. 39 is a modified version of the flowchart of FIG. 34, showing how to copy a  
20 snapshot delta from a snapshot copy facility that keeps track of blocks of a production  
21 file system that are not in use in the snapshot copies of the production file system;

1           FIG. 40 shows a specific instance of the flowchart of FIG. 39 for a snapshot copy  
2 facility that uses respective meta bit maps for indicating the blocks in a production file  
3 system that are not used in the snapshot copies of the production file system;

4           FIG. 41 is a flowchart for writing a specified data block to the production file  
5 system for a snapshot copy facility that uses a meta bit map for indicating the blocks in a  
6 production file system that are not presently in use;

7           FIG. 42 is a diagram of a preferred organization of multiple snapshots in the  
8 snapshot copy facility;

9           FIG. 43 shows a bit map including a page table and a set of pages, for use in the  
10 snapshot organization of FIG. 42;

11          FIG. 44 shows a block map including a hash table and hash lists, for use in the  
12 snapshot organization of FIG. 42;

13          FIG. 45 shows a specific construction for and interpretation of a meta bit map for  
14 the production volume; and

15          FIG. 46 shows an alternative interpretation of a meta bit map for the production  
16 volume.

17          While the invention is susceptible to various modifications and alternative forms,  
18 specific embodiments thereof have been shown by way of example in the drawings and  
19 will be described in detail. It should be understood, however, that it is not intended to  
20 limit the form of the invention to the particular forms shown, but on the contrary, the  
21 intention is to cover all modifications, equivalents, and alternatives falling within the  
22 scope of the invention as defined by the appended claims.

## DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

The present invention was developed to facilitate data recovery in a data network using a remote data replication facility. Therefore, the following detailed description will first describe the remote data replication facility, and will then describe data recovery for the remote data replication facility.

### Remote Data Replication Facility

FIG. 1 shows an IP network 220 including multiple network file servers 221, 222, and multiple hosts 223, 224, 225. The hosts and network file servers, for example, can be distributed world wide and linked via the Internet. Each of the network file servers 221, 222, for example, has multiple data movers 226, 227, 228, 232, 233, 234, for moving data between the IP network 220 and the cached disk array 229, 235, and a control station 230, 236 connected via a dedicated dual-redundant data link 231, 237 among the data movers for configuring the data movers and the cached disk array 229, 235. Further details regarding the network file servers 221, 222 are found in Vahalia et al., U.S. Patent 5,893,140, incorporated herein by reference.

In operation, it is desired for each of the network file servers 221, 222 to provide read-only access to a copy of the same file system. For example, each of the network file servers could be programmed to respond to user requests to access the same Internet site. The IP network 220 routes user requests to the network file servers 221, 222 in the same continent or geographic region as the user. In this fashion, the user load is shared among the network file servers.

1 In the wide-area network of FIG. 1, it is desired to perform read-write updating of  
2 the respective file system copies in the network file servers 221, 222 while permitting  
3 concurrent read-only access by the hosts. It is also desired to distribute the updates over  
4 the IP network.

5 There are a number of ways that updates could be distributed over the IP network  
6 from a primary data mover to multiple secondary data movers. As shown in FIG. 2, for  
7 example, a primary data mover establishes a connection 242, 243, 244 in accordance with  
8 the industry-standard Transmission Control Protocol (TCP) over the IP network 220 to  
9 each secondary data mover 245, 246, 247, and then concurrently sends the updates to  
10 each secondary data mover over the TCP connection. When the updates need to be  
11 distributed to a large number of secondary data movers, however, the amount of time for  
12 distributing the updates may become excessive due to limited resources (CPU execution  
13 cycles, connection state, or bandwidth) of the primary data mover 241. One way of  
14 extending these limited resources would be to use existing IP routers and switches to  
15 implement "fan out" from the primary data mover 241 to the secondary data movers 245,  
16 246, 247. Still, a mechanism for reliability should be layered over the Internet Protocol.

17 FIG. 3 shows that the time for distributing updates from a primary data mover 251  
18 to a large number of secondary data movers 254, 255, 256, 257 can be reduced by using  
19 intermediate data movers 252, 253 as forwarders. The primary data mover 251 sends the  
20 updates to the forwarder data movers 252, 253, and each of the forwarder data movers  
21 sends the updates to a respective number of secondary data movers. The forwarder data  
22 movers 252, 253 may themselves be secondary data movers; in other words, each may  
23 apply the updates to its own copy of the replicated read-only file system. The distribution

1 from the primary data mover 251 to the forwarder data movers 252, 253 can be done in a  
2 fashion suitable for wide-area distribution (such as over TCP connections). The  
3 forwarding method of replication of FIG. 3 also has the advantage that the distribution  
4 from each forwarder data mover to its respective data movers can be done in a different  
5 way most suitable for a local area or region of the network. For example, some of the  
6 forwarder data movers could use TCP connections, and others could use a combination of  
7 TCP connections for control and UDP for data transmission, and still other forwarders  
8 could be connected to their secondary data movers by a dedicated local area network.

9 For implementing the replication method of FIG. 3 over the Internet Protocol,  
10 there are a number of desired attributes. It is desired to maintain independence between  
11 the primary data mover and each of the secondary data movers. For example, a new  
12 secondary data mover can be added at any time to replicate an additional remote copy.  
13 The primary data mover should continue to function even if a secondary data mover  
14 becomes inoperative. It is also desired to maintain independence between the replication  
15 method and the IP transport method. Replication should continue to run even if the IP  
16 transport is temporarily inactive. It is desired to recover in a consistent fashion from a  
17 panic or shutdown and reboot. A record or log of the progress of the replication can be  
18 stored for recovery after an interruption. It is desired to build re-usable program blocks  
19 for the replication function, so that the program blocks for the replication function can be  
20 used independent of the location of the primary file system or its replicas.

21 In a preferred implementation, independence between the replication process, the  
22 IP transport method, and the primary file system being replicated, is ensured by use of a  
23 save volume. The save volume is a buffer between the data producer (i.e., the host or

1 application updating the primary file system), the replication process, and the data  
2 consumer (the secondary data movers). The save volume stores the progress of the  
3 replication over the Internet Protocol so as to maintain the consistency of the replication  
4 process upon panic, reboot, and recovery. The transport process need not depend on any  
5 "in memory" replication information other than the information in the save volume, so as  
6 to permit the replication process to be started or terminated easily on any data mover for  
7 load shifting or load balancing.

8 When a save volume is used, it can be shared between a primary data mover and a  
9 secondary data mover in the case of local file system replication, or a primary copy of the  
10 shared volume can be kept at the primary site, and a secondary copy of the shared volume  
11 can be kept at the secondary site, in the case of remote file system replication.

12 For the case of local file system replication, FIG. 4 shows a primary site including  
13 a primary data mover 260 managing access to a primary file system 261, and a secondary  
14 data mover 262 managing access to a secondary file system 263 maintained as a read-  
15 only copy of the primary file system 261. A save volume 264 is shared between the  
16 primary data mover 260 and the secondary data mover 262. This sharing is practical  
17 when the secondary site is relatively close to the primary site. A redo log 265 records a  
18 log of modifications to the primary file system 261 during the replication process for  
19 additional protection from an interruption that would require a reboot and recovery.

20 Local replication can be used to replicate files within the same network file  
21 server. For example, in the network file server 221 in FIG. 1, the primary data mover  
22 could be the data mover 226, the secondary data mover could be the data mover 227, the

1 save volume could be stored in the cached disk array 229, and replication control  
2 messages could be transmitted between the data movers over the data link 231.

3 For the case of remote file system replication, FIG. 5 shows a primary site  
4 including a primary data mover 270 managing access to a primary file system 271, and a  
5 secondary data mover 272 managing access to a secondary file system 273 maintained as  
6 a read-only copy of the primary file system 271. The primary site includes a primary  
7 save volume 274, and the remote site includes a secondary save volume 275. A redo log  
8 276 records a log of modifications to the primary file system 271 during the replication  
9 process for additional protection from an interruption that would require a reboot and  
10 recovery.

11 FIG. 6 shows a method of operating the system of FIG. 4 for local replication. In  
12 a first step 281, the primary data mover migrates a copy of the primary file system to  
13 create a secondary file system at the secondary site in such a way to permit concurrent  
14 write access to the primary file system. The migration, for example, may use the method  
15 shown in FIG. 17 of Ofek U.S. Patent 5,901,327, in which a bit map indicates remote  
16 write pending blocks. Alternatively, the migration may use a snapshot copy mechanism,  
17 for example, as described in Kedem, U.S. Patent 6,076,148, in which a bit map indicates  
18 the blocks that have changed since the time of snap-shotting of the primary file system.  
19 The snapshot method is preferred, because it is most compatible with the delta set  
20 technique for remote copy of subsequent modifications. For example, a snapshot  
21 manager creates a snapshot copy of the primary file system, as will be further described  
22 below with reference to FIGS. 25 to 27. The migration may involve taking a first  
23 snapshot copy of the primary file system, copying the first snapshot copy to the

1 secondary file system, starting replication and creating a second snapshot copy of the  
2 primary file system, and then copying the differential between the first snapshot copy and  
3 the second snapshot copy from the primary file system to the secondary file system, in a  
4 fashion similar to that described further below with reference to steps 536 and 537 in  
5 FIG. 28. In any event, it is desired for the secondary file system to become a copy of the  
6 state of the primary file system existing at some point of time, with any subsequent  
7 modifications of the primary file system being transferred through the shared save  
8 volume.

9 In step 282, the primary data mover writes subsequent modifications of the  
10 primary file system to the shared save volume. In step 283, the secondary data mover  
11 reads the subsequent modifications from the shared save volume and writes them to the  
12 secondary file system. In step 284, the secondary data mover provides user read-only  
13 access to consistent views of the secondary file system, by integrating the subsequent  
14 revisions into the secondary file system while providing concurrent read-only access to  
15 the secondary file system. This can be done by using a remote data facility to maintain a  
16 remote copy of a pair of delta volumes and to switch between the delta volumes as  
17 described in Suchitra Raman, et al., U.S. Patent Application Ser. No. 10/147,751 filed  
18 May 16, 2002, entitled "Replication of Remote Copy Data for Internet Protocol (IP)  
19 transmission, incorporated herein by reference. This can also be done by successively  
20 accumulating delta sets at the secondary, and accessing a block index or map of updated  
21 data blocks in the delta set.

22 Each delta set, for example, is a set of consistent updates to blocks of the  
23 secondary file system. The updated blocks are included in "delta chunks" transmitted

1 from the primary file server to the secondary file server. Each delta set includes a  
2 number of delta chunks, and each delta chunk belongs to a single delta set. Each delta  
3 chunk is transmitted with a header including a generation count, a delta set number, and a  
4 chunk number within the delta set. When a delta set is accumulated, a block index or  
5 map is created to identify the updated blocks in the delta set. Once the delta set has been  
6 accumulated, a background integration process is started that writes the updates into the  
7 secondary file system, and the progress of the background process is recorded by  
8 updating the block index or map for the delta set when each updated block is written to  
9 the secondary file system. During the background process, read access to the secondary  
10 file system is permitted on a priority basis. For this read access to a specified block of  
11 the file system, the secondary file server first accesses the block index or map of the delta  
12 set, and if the specified block is in the delta set, the secondary file server returns the data  
13 of the block from the delta set. If the specified block is not in the delta set, then the  
14 secondary file server returns the data of the block from the secondary file system.

15 In FIG. 6, execution loops from step 284 back to step 282. In this fashion, the  
16 secondary file system is updated from the primary site concurrently with read-only access  
17 at the secondary site.

18 FIG. 7 shows a method of operating the system of FIG. 5 for remote replication.  
19 In a first step 291, the primary data mover migrates a copy of the primary file system to  
20 create a secondary file system at the secondary site, in a fashion similar to step 281 in  
21 FIG. 6. In step 292, the primary data mover writes subsequent modifications of the  
22 primary file system to the primary save volume, in a fashion similar to step 282 in FIG. 6.  
23 In step 293, the modifications are copied from the primary save volume to the secondary

1 save volume, by transmitting delta chunks. In step 294, the secondary data mover reads  
2 the modifications from the secondary save volume and writes them to the secondary file  
3 system. In step 295, the secondary data mover provides user read-only access to  
4 consistent views of the secondary file system, in a fashion similar to step 284 of FIG. 6.  
5 Execution loops from step 295 back to step 292. In this fashion, the secondary file  
6 system is remotely updated from the primary site concurrently with read-only access at  
7 the secondary site.

8 FIG. 8 shows layered programming 300 for a primary data mover. It is desired to  
9 use layered programming in accordance with the International Standard Organization's  
10 Open Systems Interconnection (ISO/OSI) model for networking protocols and distributed  
11 applications. As is well known in the art, this OSI model defines seven network layers,  
12 namely, the physical layer, the data link layer, the network layer, the transport layer, the  
13 session layer, the presentation layer, and the application layer.

14 As shown in FIG. 8, the layered programming 300 includes a conventional  
15 TCP/IP transport layer 301. The layers above the TCP/IP transport layer 301 include a  
16 replication control protocol (RCP) session layer 302, a volume multicast presentation  
17 layer 303, and an IP-FS (file system) copy send-thread 304 and an IP-replication send-  
18 thread 305 at the program layer level. Over these program layers is a management and  
19 configuration command interpreter (MAC\_CMD) 306 for system operator set-up,  
20 initiation, and supervisory control of the replication process.

21 In operation, the RCP layer 302 provides an application program interface (API)  
22 for multicasting data over TCP/IP. RCP provides callback, acknowledgement (ACK),  
23 and resumption of aborted transfers.

1 RCP provides the capability for a remote site to replicate and rebroadcast remote  
2 copy data. The remote site functions as a router when it rebroadcasts the remote copy  
3 data. RCP can also be used to replicate data locally within a group of data movers that  
4 share a data storage system.

5 To create a new remote copy in response to a supervisory command, the  
6 command interpreter 306 initiates execution of a replication module 310 if the replication  
7 module is not presently in an active mode. Then, the command interpreter 306 invokes a  
8 snapshot manager 308 to create a snapshot copy 309 of a primary file system volume  
9 307. When the snapshot copy is created, the snapshot manager 308 obtains a current  
10 delta set number from the replication module 310 and inserts the current delta set number  
11 into the metadata of the snapshot. The current delta set number for the snapshot is all that  
12 the secondary needs to identify modifications that are made subsequent to the creation of  
13 the snapshot. In this fashion, any number of new remote copies can be created at various  
14 times during operation of the replication module, with the snapshot process operating  
15 concurrently and virtually independent of the replication module. For example,  
16 whenever synchronization of a remote copy is lost, for example due to a prolonged  
17 disruption of network traffic from the primary site to the remote site, a new remote copy  
18 can be created to replace the unsynchronized remote copy.

19 Once the snapshot copy 309 is accessible, the command interpreter 306 initiates  
20 execution of an instance of the IP-FS copy send-thread 304. The instance of the IP-FS  
21 copy send-thread 304 reads data from the snapshot copy 309 and calls upon the volume  
22 multicast layer 303 to multicast the remote copy data to all of the secondary data movers  
23 where the remote copies are to be created. This can be a copy by extent, so there is no

1 copying of invalid or unused data blocks. For example, the volume multicast layer 303 is  
2 given a copy command (@vol., length) specifying a volume and an extent to be copied,  
3 and may also specify a group of destinations (an RCP group). The snapshot copy 309 of  
4 the primary file system identifies the next valid block to be copied, and the number of  
5 valid contiguous blocks following the next block. These blocks are copied at the logical  
6 level, so it does not matter what physical structure is used for storing the secondary file  
7 system at the secondary site. The copying is done locally, or by remote copy, for  
8 example by transporting the data block over IP. The volume multicast layer 303 invokes  
9 the RCP layer 302 to transport each data block.

10 During the remote copy process, whenever a modification is made to a block of  
11 the primary file system volume 307, the replication module 310 logs an indication of the  
12 modified block in a log 314 and later assembles the modification into a delta set chunk  
13 written to a primary save volume 311. The replication module 310 logs the indications in  
14 the log 314 on a priority or foreground basis as data is written to the primary file system  
15 volume 307, and also logs boundaries between delta sets. The replication module 310  
16 later reads the log 314 to read the indicated modifications from the primary file system  
17 volume 307, assemble the indicated modifications into delta set chunks on a background  
18 basis, and store the delta set chunks in a save volume chunk area of the save volume 311.  
19 For example, the log is in the form of a queue of two bit-map tables, a new one of the  
20 tables being written to coincident with write operations upon the primary file system  
21 volume 307, and an old one of the tables being read to determine blocks to copy from the  
22 primary file system to create a new delta set in the save volume 311. When the delta set  
23 chunks become available for distribution from the save volume 311, the replication

1 module 310 updates the save volume mailbox area 312 by storing each delta set chunk  
2 definition (@vol., length).

3 The IP-replication send-thread instance 305 polls the save volume mailbox area  
4 312 to see if any delta set chunks have been stored in the save volume chunk area 313. If  
5 so, then the thread instance calls upon the volume multicast layer 303 to multicast the  
6 delta set chunks to the data movers that manage the storage of the respective remote file  
7 system copies. For example, for each delta set chunk, the IP-replication send-thread  
8 instance 305 issues a volume multicast command to the volume multicast layer 303.  
9 When the chunk multicast is completed, the IP-replication send-thread instance 305  
10 updates its context on the save volume 311 in the mailbox area 312. At reboot after an  
11 interruption of multicast of a chunk, the IP-replication send-thread instance is able to  
12 restart the multicast of the chunk. The IP-replication send-thread instance also is  
13 responsible for retrying transmission of the chunk whenever the connection with the  
14 secondary is interrupted.

15 FIG. 9 shows the layered programming 320 for a secondary data mover. The  
16 programming includes a TCP/IP layer 321, an RCP layer 322, a volume multicast layer  
17 323, and a management and configuration command interpreter (MAC\_CMD) 324.  
18 During creation of a new remote copy in a secondary file system volume 325, the volume  
19 multicast layer 323 writes remote copy data from the primary data mover to the  
20 secondary file system volume 325, and concurrently writes modifications (delta set  
21 chunks) from the primary data mover to a save volume chunk area 326 of a secondary  
22 save volume 327.

1       A header for the changes in a next version of the delta set is sent last, because  
2       there is no guarantee of the order of receipt of the IP packets. The header of the delta set  
3       includes a generation count, the number of delta blocks for the next version of the delta  
4       set, a checksum for the header, and a checksum for the data of all the delta blocks. The  
5       receiver checks whether all of the changes indicated in the header have been received.

6       Once a complete remote snapshot copy has been reconstructed in the secondary  
7       file system volume 325, a playback module 328 is activated to read the modifications  
8       from the save volume chunk area 326 and integrates them into the secondary file system  
9       volume 325. From each delta-set chunk in the save volume area 326, the playback  
10      module 328 gets the block address and number of contiguous blocks to be written to the  
11      secondary file system volume. An access module 329 provides read-only access to a  
12      consistent view of the secondary file system in the secondary file system volume 325.

13      FIG. 10 shows a procedure executed by the primary site of FIG. 8 to perform  
14      replication of the primary file system. When replication is started in a first step 341, the  
15      primary file system is paused to make it consistent. Migration of the primary file system  
16      to the secondaries can then be started using a remote copy facility or snapshot manager.  
17      Then, in step 342, concurrent write access to the primary file system is resumed, and all  
18      modifications made on the primary file system are logged at the volume level on a  
19      priority or foreground basis when each modification is made. In addition, a background  
20      process of delta-set creation is initiated.

21      Two configurable triggers specify the rate of delta set creation: a timeout  
22      parameter and a high water mark parameter. Whenever delta set creation is initiated, the  
23      current time, as indicated by a real-time clock, is added to a configurable timeout interval

1 to produce the timeout parameter. The high water mark specifies an amount of modified  
2 data, in megabytes. The first trigger that occurs will trigger the creation of a delta set.  
3 The replication module creates the delta set by pausing the primary file system, copying  
4 the modified blocks from the primary file system to the delta set volume, and then  
5 resuming the primary file system. By logging indications of the modified blocks and  
6 later copying the modified blocks, multiple modifications to the same block are  
7 represented and transported once during a single delta set.

8 In step 343, the background process of delta set creation is temporarily suspended,  
9 for example, by placing the process on a task queue that is periodically serviced. In step  
10 344, execution of the delta set creation process is resumed. In step 345, the modification  
11 size is compared to the high water mark. If the high water mark is not exceeded, then  
12 execution continues to step 346. In step 346, the present value of the real-time clock is  
13 compared to the timeout parameter. If the timeout parameter has not been exceeded, then  
14 execution loops back to step 343. Otherwise, execution continues to step 347. Execution  
15 also branches to step 347 from step 345 if the modification size is greater than the high  
16 water mark.

17 In step 347, the primary file system is paused. In step 348, a new delta set is  
18 created by starting the copying of modified blocks from the primary file system volume  
19 to the new delta set. In step 349, the logging of new modifications into a new table is  
20 started. In step 350, the time-out and high water mark is re-armed. In other words, a new  
21 value for the timeout parameter is computed as the current real time plus the configurable  
22 timeout interval, and the modification size is reset to indicate the size of the new

1 modifications. In step 351, the primary file system is resumed. Execution loops from  
2 step 351 back to step 343 to suspend the background process of delta set creation.

3 To maintain the consistency of the delta set created in the primary save volume,  
4 the primary file system could remain paused and not resumed in step 351 until the copy  
5 process begun in step 348 is completed. Preferably, however, the copy process begun in  
6 step 348 is a snapshot copy process, so that write access to the primary file system may  
7 resume in step 351 before the copy process has been completed. For the example of the  
8 modification log being a queue of two bit-map tables, when a write access to a block in  
9 the primary file system is requested, the old bit map is accessed on a priority basis. If the  
10 corresponding bit in the old bit map indicates a modified block in the primary file system  
11 volume not yet copied to the save volume, then it is copied on a priority basis to the save  
12 volume before the new write data is written to the primary file system volume. As soon  
13 as a modified block has been copied from the primary file system volume to the save  
14 volume, the corresponding bit in the old bit map is cleared. In this fashion, at the  
15 completion of the copy process, the entire old table will be in a reset state, ready to be  
16 used as the next new table.

17 When the copy process started in step 348 is completed, the replication module  
18 sets the save volume mailbox area to show that a new delta set is ready for transmission.  
19 Upon polling the mailbox area, the IP-replication send-thread finds that the new delta set  
20 is ready for transmission, and invokes the volume multicast layer to transmit the delta set  
21 to the secondary sites. After step 351, execution loops back to step 343.

22 FIG. 11 shows a flow chart of the overall procedure of creating a new remote  
23 copy, either for the first time at a secondary site or as a replacement for a remote copy

1 that needs to be resynchronized with the primary file system. In a first step 352, the  
2 snapshot manager creates a snapshot copy of the primary file system at the end of any  
3 pending transaction upon the primary file system (e.g., when the primary file system  
4 becomes consistent after it is paused in step 341 of FIG. 10 or in step 347 of FIG. 10.)  
5 The replication module independently writes any subsequent modifications into a current  
6 delta set for the next transaction.

7 In step 353, the snapshot manager obtains the current delta set number from the  
8 replication module and inserts it into metadata of the snapshot copy. In step 354, the IP-  
9 FS copy send-thread is started in order to send volume extents of the snapshot copy to the  
10 secondary data mover, by invoking the volume multicast layer for each extent.

11 In step 355, when the IP-FS copy send-thread is finished, the primary data mover  
12 sends a "start playback" signal to the secondary data mover. In step 356, the secondary  
13 data mover receives the "start playback" signal from the primary data mover, and starts  
14 the playback module. In step 357, playback module begins playback from the delta set  
15 indicated by the delta set number in the snapshot metadata.

16 The playback module (328 in FIG. 23) at the secondary site integrates the delta  
17 set modifications into secondary file system. Each time that a new delta set appears in  
18 the secondary save volume, the modifications can be integrated into the secondary file  
19 system, for example, by pausing the secondary file system, copying the modifications  
20 from the secondary save volume into the secondary file system, and resuming the  
21 secondary file system. Alternatively, a timeout interval and a high water mark value can  
22 be configured for the secondary site, so that the modifications may be integrated into the  
23 secondary file system at a rate less frequent than the rate at which the new delta sets

1 appear in the secondary save volume. In this case, the modifications from the secondary  
2 save volume would not be integrated into the secondary file system until the timeout time  
3 is reached unless the amount of modifications in the save volume reaches the high water  
4 mark. As described above, the integration of the modifications can be performed  
5 concurrently with read-only access to a consistent view of the secondary file system.

6 FIG. 12 shows a flowchart of the IP-replication send-thread (305 in FIG. 8). In a  
7 first step 361, the thread polls the primary save volume mailbox area. If the mailbox area  
8 indicates that there is not a new delta set chunk in the primary save volume area, then the  
9 thread is finished for the present task invocation interval. Execution of the thread is  
10 suspended in step 363, and resumed in step 364 at the next task invocation interval.

11 If the mailbox area indicates that there is a new delta set chunk in the primary  
12 save volume, then execution continues from step 362 to step 365. In step 365, the IP-  
13 replication send-thread issues a volume multicast command to broadcast or forward the  
14 delta set chunk to specified destination data movers. In step 366, if the multicast has  
15 been successful, then execution branches to step 367. In step 367, the IP-replication  
16 send-thread updates the primary save volume mailbox to indicate completion of the  
17 multicast, and execution continues to step 363 to suspend execution of the thread until the  
18 next task invocation interval.

19 In step 366, if the multicast is not successful, then execution continues to step 368  
20 to test whether more than a certain number (N) of retries have been attempted. If not,  
21 then execution loops back to step 365 to retry the multicast of step 365. If more than N  
22 retries have been attempted, then execution continues from step 368 to step 369. In step

369, the IP-replication send-thread logs the error, and then in step 370, passes execution to an error handler.

FIG. 13 shows various objects defined by the volume multicast layer. The volume multicast layer provides multicast service to instances of a VolMCast object 370 representing a volume multicast relationship between a respective primary file system volume specified by a volume name (volumeName) and a respective group of secondary data movers specified by an RCP group name (rcpgpeName). For example, at configuration time, one or more RCP groups are defined in response to configuration commands such as:

```
.RCP_config <server_name> add <IP>
```

This configuration command adds the IP address (IP) of a specified destination data mover (server\_name) to an RCP group.

Also at configuration time, a specified data mover can be defined to be a primary data mover with respect to the RCP group (a relationship called a MultiCastNode) in response to a configuration command such as:

```
.server_config <server_name> rep <groupname> add <IP>
```

where “server\_name” is the name for the primary data mover, “groupname” is the name of a configured RCP group, and “IP” is the IP address of the primary data mover. When configuration of the MultiCastNode object is finished, the MultiCastNode object will

1 have its own name, a name for the primary data mover, an RCP group name, and a list of  
2 IP addresses to which the primary server should broadcast in order to transmit IP packets  
3 to all the secondary data movers in the RCP group.

4 The VolMCast object can then be built on top of a MultiCastNode object. The  
5 additional information required for the VolMCast object is, on the sender side, the  
6 primary or source file system volume and on each receiver side, the secondary or  
7 destination file system volume. For flexibility, it is permitted to specify a different  
8 volume name on each secondary data mover. By specifying the destination volume  
9 names during creation of the VolMCast object, it is not necessary to specify the  
10 destination volume names at each copy time. For example, the VolMCast object is  
11 defined by configuration commands to the primary data mover such as:

```
12  
13 .server_config <server_name> "volmcast MultiCastNodeName>  
14 [ -src | -dest ] volume"
```

15  
16 where <server\_name> is the name of the MultiCast Node.

17 Once the VolMCast object has been defined, an IP-replication service can be  
18 configured for the object upon the primary data mover. Then the primary data mover will  
19 respond to commands for starting the replication service and stopping the replication  
20 service upon the VolMCast object. When replication is stopped on a secondary, the  
21 secondary file system is left in a consistent state. In other words, if a replay was in  
22 progress, the stop will complete when the replay is finished.

1       The primary data mover may respond to additional commands for create a new  
2 delta set on demand, updating the replication policy (high water mark and timeout  
3 interval parameters) on the primary file system or secondary file systems, and defining  
4 persistency of the replication process upon remount or reboot of the primary file system  
5 or any one of the secondary file systems. For example, at reboot the replication service is  
6 re-started on the primary file system and the secondary file system in the state it was at  
7 unmount or shutdown. A recovery of the replication context happens at reboot or on  
8 remount. The replica recovery is executed before the primary and secondary file systems  
9 are made available for user access. This allows all modifications during the recovery of  
10 the primary file system to be logged by the replication service.

11       As shown in FIG. 13, the volume multicast layer is responsive to a number of  
12 commands 371 from higher layers in the protocol stack. In addition to the configuration  
13 commands for defining a new VolMCast object relating a specified primary file system  
14 volume to a specified RCP group, an existing VolMCast object can be opened for either a  
15 sender mode or a receiver mode. An opened VolMCast object can be closed. Once a  
16 VolMCast object has been opened in a sender mode, it can be called upon to broadcast a  
17 control block (CB) to the secondary volumes of the VolMCast object, such as a control  
18 block specifying a remote copy of a specified extent of the primary volume.

19       Control blocks may specify various operations upon the secondary volumes of the  
20 VolMCast object, such as cluster file system commands for performing operations such  
21 as invalidations, deletions, renaming, or other changes in the configuration of the objects  
22 of the file system upon all copies (local or remote) of the file system. In this case, RCP is  
23 used for the broadcast or forwarding of the cluster file system commands to all the data

1 movers that are to operate upon the local or remote copies of the file system, and for  
2 returning acknowledgement of completion of the operations upon all of the copies of the  
3 file system.

4 With reference to FIG. 13, the volume multicast layer defines a VolMCastSender  
5 object 372 instantiated when a VolMCast instance is opened in the sending mode, and a  
6 VolMCastReceiver object 373 instantiated when a VolMCast instance is opened in a  
7 receiving mode. The VolMCastSender object class and the VolMCastReceiver object  
8 class inherit properties of the VolMCast object class. When the volume multicast layer is  
9 called upon in a primary data mover to maintain remote copies of a specified extent of a  
10 VolMCastSender instance, an instance of a VolMCastCopy thread 374 is created and  
11 executed. The VolMCastCopy thread instance accesses the delta sets from a primary  
12 save volume 375 to produce a write stream 376 of blocks sent down to the RCP layer. At  
13 the secondary data mover, an instance of a VolMCastReceiver thread 377 is instantiated  
14 and executed to receive a read stream 378 of blocks and write the copied delta sets into a  
15 secondary save volume 379. An instance of an acknowledgement thread 380 returns an  
16 acknowledgement 381 of completion of copying of a delta-set for an extent to the  
17 secondary file system. The acknowledgement is sent down to the RCP layer of the  
18 secondary data mover. At the primary, the RCP layer sends the acknowledgement 382 to  
19 an instance of an acknowledgement thread 383.

20 RCP is a session-layer protocol, for replication from one primary to multiple  
21 secondary sites. Control is initiated by the primary, except when recovering from aborted  
22 transfers. RCP uses TCP between the primary and secondary for control and data.

1 Network distribution is by an application-level multicast (ALM) using the RCP as a  
2 forwarder. Port sharing with HTTP is used for crossing firewalls.

3 RCP may support other replication applications in addition to 1-to-N IP-based  
4 replication for wide-area distribution of read-only data. These other applications include  
5 1-to-N volume mirroring, cluster file system commands, remote file system replication,  
6 and distribution and replication of other commands that may be recognized by the data  
7 movers.

8 The 1-to-N volume mirroring is a simplification of to 1-to-N IP-based replication  
9 for wide-area distribution of read-only data, because the volume mirroring need not  
10 synchronize a remote volume with any consistent version of the primary volume until the  
11 remote volume needs to be accessed for recovery purposes.

12 Remote file system replication also uses RCP for broadcast or forwarding an  
13 application command to a remote data mover to initiate a replication of a file system  
14 managed by the remote data mover. In a similar fashion, RCP may broadcast or forward  
15 other commands recognized by data movers, such as iSCSI or remote-control type  
16 commands for archival storage. For example, RCP could broadcast or forward remote  
17 control commands of the kind described in Dunham, U.S. Patent 6,353,878 issued March  
18 5, 2002 entitled "Remote Control of Backup Media in a Secondary Storage Subsystem  
19 Through Access to a Primary Storage Subsystem," incorporated herein by reference.

20 The RCP forwarder is composed of two RCP sessions, an outbound session at the  
21 primary, and an inbound session at the secondary. The inbound RCP session receives a  
22 group name and looks up the group in a routing table. If routes for the group exist in the

1 routing table, then an RCP forwarder is created at the secondary, including a data path by  
2 pointer passing from an “in” session to an “out” session.

3 An RCP group may be configured to include application-level multicast (ALM)  
4 topology. For example, ALM route configuration commands begin with an identifier  
5 number for the network file server (“cel”) that contains the forwarder data mover, and an  
6 identifier number (“ser”) for the forwarder data mover in the network server. The  
7 configuration commands end with a “nexthop” specification of an immediate destination  
8 data mover:

9  
10 cell1-ser2: rcproute add group=g1 nexthop=cell2-ser2

11 cel2-ser2: rcproute add group=g1 nexthop=cell2-ser3

12 cel2-ser2: rcproute add group=g1 nexthop=cell2-ser4

13  
14 In effect, the forwarder data mover adds the “nexthop” specification to an entry  
15 for the RCP group in the routing table in the forwarder data mover. This entire entry can  
16 be displayed by the following configuration command:

17  
18 cel2-ser2: rcproute display

19  
20 The entry is displayed, for example, as a list of the “nexthop” destination data  
21 movers. The entry can be deleted by the following configuration command:

22  
23 cel2-ser2: rcproute delete

1 Each immediate destination data mover may itself be configured as a forwarder in  
2 the RCP group. In this case, RCP commands and data will be forwarded more than once,  
3 through a chain of forwarders. The set of possible RCP routes from a primary or  
4 forwarder in effect becomes a tree or hierarchy of destinations.

5 The ALM commands may also include commands for creating sessions and  
6 sending control blocks or data. For example, the following ALM command creates a  
7 session and sends application data to all destinations in group "g1" from cell1-ser2 from a  
8 file (named "filename") using a test application (named "rcpfiletest").

9  
10 cell1-ser2: rcpfiletest data=filename group=g1  
11

12 FIG. 14 shows the RCP collector service 390 at a primary site. The programming  
13 for the RCP collector service includes an RCP session manager 391, collector and worker  
14 threads 392, and a single-thread RCP daemon 393. The RCP session manager 391  
15 responds to requests from higher levels in the protocol stack, such as a request from an  
16 application 394 to open an RCP pipe 395 between the application 394 and the RCP  
17 collector service 390. The application 394 may then send to the session manager 391  
18 requests to setup sessions with RCP groups. A session queue 396 stores the state of each  
19 session, and a control block queue 397 keeps track of control blocks sent via TCP/IP to  
20 the secondary data movers in the RCP groups. An RCP routing table 398 identifies the  
21 immediate destinations of each RCP group to which the TCP/IP messages from the RCP  
22 collection service are to be sent, as well as any other destinations to which the messages  
23 will be forwarded. For communication of the TCP/IP messages between the RCP service

1 and the network, TCP port :80 is opened in both directions (i.e., for input and output).

2 The single thread RCP daemon 393 is used for interfacing with this TCP port :80.

3 FIG. 15 shows the RCP collector service 400 at a secondary site. The RCP  
4 collector service at the secondary site is similar to the RCP collector service at the  
5 primary site, in that it includes an RCP session manager 401, collector and worker  
6 threads 402, a single thread RCP daemon 403 for access to/from TCP port :80, an RCP  
7 session state queue 406, an RCP control block queue 407, and an RCP routing table 408.  
8 The primary difference between the RCP collector service at the secondary site from the  
9 RCP collector service at the primary site is in the collector and worker threads 402. At  
10 the RCP secondary, the RCP commands and data are received from the TCP port :80  
11 instead of from the application 404. The application 404 is the consumer of the RCP  
12 data, instead of a source for RCP data. The RCP collector service 400 at the secondary  
13 site may also serve as a forwarder for RCP commands, and therefore the RCP collector  
14 service and worker threads 402 at the secondary site include a forwarder thread that does  
15 not have a similar or complementary thread in the RCP collector service at the primary  
16 site.

17 In operation, an application 404 can initialize the RCP collector service so that the  
18 RCP collector service will call back the application upon receipt of certain RCP  
19 commands from TCP port :80. For example, if a new connection command is received  
20 from TCP port :80, then the RCP daemon 403 forwards the new connection command to  
21 the RCP session manager. The RCP session manager 401 recognizes that this connection  
22 command is associated with an application 404 at the secondary site, opens an RCP pipe  
23 405 to this application, and calls the application 404 indicating that the RCP pipe 405 has

1    been opened for the RCP session. (The volume multicast receiver thread 377 of FIG. 13  
2    is an example of such an application.) The application 404 returns an acknowledgement.  
3    If the new connection is for a new RCP session, then the session manager creates a new  
4    RCP session, and places state information for the new session on the RCP session queue  
5    406. RCP control blocks and data may be received for the session from the TCP port :80.  
6    The data may be forwarded to the application, or to a file specified by the application.  
7    RCP control blocks to be executed by the RCP collector service 400 may be temporarily  
8    placed on the control block queue 407. RCP control blocks or data intended for other  
9    secondary site may be forwarded to the intended secondary sites.

10       FIG. 16 shows further details of the forwarding of RCP commands and data by a  
11    data mover 430 identified as Cel2-Ser1. The data mover 430 is programmed with a  
12    TCP/IP layer 431 for communication with the IP network 220, and an RCP layer 432  
13    over the TCP/IP layer. For forwarding the RCP commands and data, the RCP layer 432  
14    creates an inbound session 433 and an outbound session 434. The inbound session 433  
15    receives RCP commands from the TCP/IP layer 431. The TCP/IP data stream is retained  
16    in a data buffer 435. When an RCP command calls for the forwarding of RCP commands  
17    or data to another data mover in a specified RCP group, the inbound session 433  
18    performs a lookup for the group in a routing table 436.

19       In the example of FIG. 16, the routing table 436 includes a copy of all of the  
20    routing information for each group of which the data mover 430 is a member. In this  
21    case, for GROUP1, the primary data mover sends RCP commands to at least data movers  
22    CEL2-SER1 and CEL9-SER1. CEL2-SER1 (i.e., the data mover 430) forwards the RCP  
23    commands and RCP data to data movers CEL3-SER1 and CEL7-SER1. In particular, the

1 inbound session 433 creates an outbound session 434 and creates a TCP/IP data path  
2 from the inbound session 433 to the outbound session 434 by passing pointers to the data  
3 in the data buffer. The outbound session 434 invokes the RCP/IP layer 431 to multicast  
4 the TCP data stream in the data buffer 435 over the IP network 220 to the data movers  
5 CEL3-SER1 and CEL7-SER1.

6 The data mover CEL3-SER1 in succession forwards the RCP commands to data  
7 movers CEL4-SER1 and CEL5-SER1. Normally, the data mover CEL2-SER1 (430)  
8 does not need to know that the data mover CEL3-SER1 forwards the RCP commands to  
9 data movers CEL4-SER1 and CEL5-SER1, but if the data mover CEL2-SER1 (430)  
10 would fail to receive an acknowledgement from CEL3-SER1, then the data mover CEL2-  
11 SER1 could minimize the impact of a failure of CEL3-SER1 by forwarding the RCP  
12 commands to CEL4-SER1 and CEL5-SER1 until the failure of CEL3-SER1 could be  
13 corrected.

14 FIG. 17 shows a flowchart of how the RCP collector service at the secondary site  
15 processes an inbound RCP session command. In a first step 411, the RCP collector  
16 service receives a session command. In step 412, if this session command is not a  
17 command to be forwarded to other secondary sites, then execution branches to step 413  
18 to execute the action of the command, and the processing of the session command is  
19 finished.

20 In step 412, if the session command is a command to be forwarded to other  
21 secondary sites, then execution continues from step 412 to step 414. In step 414, the  
22 RCP collector service gets the RCP group name from the session command. Then, in  
23 step 415, the RCP collector service looks up the group name in the RCP routing table

1 (408 in FIG. 15). If the group name is not found, then execution branches from step 416  
2 to step 417. In step 417, the RCP collector service returns an error message to the sender  
3 of the session command.

4 In step 416, if the group name is found in the RCP routing table, then execution  
5 continues from step 416 to step 418. In step 418, the RCP collector service forwards the  
6 action of the session command to each secondary in the group that is an immediate  
7 destination of the forwarder (i.e., the data mover that is the secondary presently  
8 processing the RCP session command). This is done by instantiating local replication  
9 threads or creating outbound sessions for forwarding the action of the session command  
10 to each secondary in the group that is an immediate destination of the forwarder. After  
11 step 418, processing of the RCP session command is finished.

12 FIG. 18 shows an example of forwarding and local replication. In this example,  
13 the IP network 220 connects a primary data mover 421 to a network file server 422 and a  
14 secondary data mover 423. The network file server 422 includes three data movers 424,  
15 425, and 426, and storage 427. The primary data mover manages network access to a  
16 primary file system 428. The data mover 424 functions as a forwarder data mover. The  
17 data mover 425 functions as a secondary data mover managing access from the network  
18 to a secondary file system (copy A) 429. The data mover 426 functions as a secondary  
19 data mover managing access from the network to a secondary file system (copy B) 430.  
20 The data mover 423 manages network access to a secondary file system (copy C) 431.

21 In operation, when the primary data mover 421 updates the primary file system  
22 428, it multicasts the modified logical blocks of the file system volume over the IP  
23 network 220 to the forwarder data mover 424 and to the secondary data mover 423. The

1 forwarder data mover 424 receives the modified blocks, and performs a local replication  
2 of the blocks to cause the secondary data mover 425 to update the secondary file system  
3 (copy A) 429 and the to cause the secondary data mover 426 to update the secondary file  
4 system (copy B) 430.

5 To perform the local replication, the forwarder data mover 424 has its volume  
6 multicast layer (323 in FIG. 9) save the modified blocks in a save volume 432 in the  
7 storage 427, and then the forwarder data mover 424 sends replication commands to the  
8 local secondary data movers 425 and 426. Each local secondary data mover 425, 426 has  
9 its playback module (328 in FIG. 9) replay the modifications from the save volume 432  
10 into its respective secondary file system copy 429, 430.

11 FIG. 19 shows the sharing of the data mover's network TCP port :80 (440)  
12 between HTTP and RCP. This configuration is used in all data movers having the RCP  
13 collector service; i.e., primary, secondary, or forwarder. The TCP data channel from  
14 TCP port :80 (440) provides an in-order byte stream interface. IP packets 444 for HTTP  
15 connections and IP packets 445 for RCP connections from the network 220 are directed  
16 to the data mover's TCP port :80 (440). The TCP port :80 (440) is opened in both  
17 directions (i.e., input and output). In the input direction, the data mover uses a level 5  
18 (L5) filter 441 for demultiplexing the IP packets for the HTTP connections from the IP  
19 packets for the RCP connections based on an initial segment of each TCP connection.  
20 The L5 filter hands the TCP connection off to either a HTTP collector service 442 or an  
21 RCP collector service 443. (The RCP collector service 443 is the collector service 390 in  
22 the RCP primary of FIG. 14 or the RCP collector service 400 in an RCP secondary of  
23 FIG. 15.) For example, if the initial segment of a TCP connection contains "HTTP/1.X",

1 then the L5 filter 441 directs the IP packets for the connection to the HTTP collector  
2 service 442. If the initial segment of the TCP connection contains "RCP/1.0", then the IP  
3 packets for the TCP connection are directed to the RCP collector service 443. (In an  
4 alternative arrangement, the connection could be split as is done in a conventional stand-  
5 alone IP switch.)

## 6 7 **Data Recovery With Internet Protocol Replication With Or Without Full** 8 **Resync**

9 A remote replication system may protect against data loss resulting from a  
10 disaster. For example, FIG. 20 shows the state of a remote replication system before a  
11 disaster. The remote replication system copies data in an asynchronous fashion from a  
12 primary data storage system at the primary site to a secondary data storage system at a  
13 secondary site. The primary and secondary data storage systems could be file servers, for  
14 example as shown in FIG. 1.

15 As shown in FIG. 20, the primary site includes a replication service 450  
16 replicating changes made to a primary file system 451, and the secondary site includes a  
17 playback service 452 writing the changes into a secondary file system 453. The  
18 replication service 450 transfers the changes in "delta sets" sent over an IP pipe 454 to  
19 the playback service 452. The disaster occurs at the primary site during this transfer. At  
20 this time, the playback service 452 is a few delta sets behind the replication service 450;  
21 for example, the playback service is playing back delta set (n-x) when the replication  
22 service 452 is replicating delta set (n).

1       Once the disaster causes the primary site to go down, the secondary site can be  
2    activated to service client read/write requests. When the primary site comes back up, it  
3    can be re-synchronized to the secondary site, in many cases without making a full copy of  
4    the file system. Then replication can be re-established as it was before the disaster. This  
5    recovery process has been designed to minimize data loss and the down time of both the  
6    primary and secondary site.

7       FIG. 21 is a flowchart of the preferred recovery process. Initially, in step 461,  
8    data is replicated from the primary file system at the primary site by sending delta sets to  
9    the secondary file system at the secondary site. Then in step 462, the primary site  
10   becomes inoperative. In response, in step 463, the secondary site is activated for  
11   read/write access to the secondary file system. This is done by a subroutine for failover  
12   with a checkpoint and without sync, as further described below with respect to FIG. 23.  
13   Eventually, in step 464, the primary site becomes operative. In response, in step 465, the  
14   primary file system is synchronized to the state of the secondary file system, as further  
15   described below with respect to FIG. 28. Then in step 466, read/write access to the  
16   primary file system and replication of data from the primary file system to the secondary  
17   file system is resumed in a failback operation, as further described below with reference  
18   to FIG. 30.

19       FIG. 22 shows the state of the remote replication system of FIG. 20 during the  
20   failover operation. The (n-x) deltasets are played back into the secondary file system  
21   453, to put the secondary file system 453 into a state called a "restart point". A snapshot  
22   copy facility 456 creates a snapshot 455 of the "restart point" in order to facilitate  
23   synchronization of the primary file system with the secondary file system once the

1 primary site becomes operative. The playback service 452 is stopped, and the secondary  
2 file system 453 is remounted as read/write. The snapshot copy facility 456 keeps a  
3 record of changes made to the secondary file system 453 after the "restart point."

4 FIG. 23 shows a flowchart of the failover operation. In step 471, if possible, the  
5 primary file system is remounted as a read-only file system. Next, in step 472, the  
6 secondary site finishes playback of the (n-x) delta sets that it had received but had not  
7 played back prior to the primary site becoming inoperative. In step 473, the secondary  
8 site creates a snapshot copy of the "restart point" of the secondary file system. This  
9 snapshot is stamped with the last delta set having been replayed into the secondary file  
10 system prior to the creation of the snapshot. In other words, an identifier of this last delta  
11 set is stored as an attribute of the snapshot. In general, whenever a snapshot is made of a  
12 file system during playback of delta sets into the file system, the snapshot is stamped with  
13 the last delta set having been replayed. In step 474, the secondary file system is  
14 remounted as a read/write file system. In step 475, the snapshot process retains and  
15 identifies all changes (delta) made to the secondary file system since the restarting point.

16 Preferably the snapshot copy facility 456 retains and identifies the changes at a  
17 logical volume level of data storage. For example, the present state of the secondary file  
18 system is stored in a "clone volume," and old versions of the logical blocks that have  
19 been changed in the clone volume are saved in a "save volume". In order to conserve  
20 storage, the logical blocks of the save volume are dynamically allocated to the old  
21 versions of the changed blocks as the changes are made to the clone volume. As shown  
22 in FIG. 24, for each logical block that has been changed in the clone volume, a block map  
23 480 identifies the logical block address ( $S_i$ ) of the old version of the block in the save

1 volume and the corresponding logical block address ( $B_i$ ) of the changed block in the  
2 clone volume.

3 FIG. 25 shows details of the preferred snapshot copy facility 456, which provides  
4 multiple snapshots 483, 503 of a production file system 481. (Two successive snapshots  
5 of the secondary file system are used, for example, in the resync subroutine further  
6 described below with reference to FIG. 28.) The content of each snapshot file system  
7 483, 503 is the state of the production file system 481 at a particular point in time when  
8 the snapshot was created. The snapshot copy facility 456 provides a hierarchy of objects  
9 in a volume layer 490 supporting the file systems in a file system layer 491. The  
10 production file system 481 is supported by read/write access to a file system volume 482.  
11 Each snapshot file system 483, 503 provides read-only access to a respective snapshot  
12 volume 484, 504.

13 Additional objects in the volume layer 490 of FIG. 25 permit the content of each  
14 snapshot file system to be maintained during concurrent read/write access to the  
15 production file system 481. The file system volume 482 is supported by a snapped  
16 volume 485 having read access to a clone volume 487 and write access to a delta volume  
17 486. The delta volume 486 has read/write access to the clone volume 487 and read/write  
18 access to a save volume 488.

19 In the organization of FIG. 25, the actual data is stored in blocks in the clone  
20 volume 487 and a respective save volume 488, 506 in storage for each snapshot. The  
21 delta volume 486 also accesses information stored in a bit map 489 and the block map  
22 480. The bit map 489 indicates which blocks in the clone volume 487 have prior  
23 versions in the save volume 488. In other words, for read-only access to the snapshot file

1 system, the bit map 489 indicates whether the delta volume should read each block from  
2 the clone volume 487 or from the save volume 488. For example, the bit map is stored in  
3 memory and it includes a bit for each block in the clone volume 487. The bit is clear to  
4 indicate that there is no prior version of the block in the save volume 488, and the bit is  
5 set to indicate that there is a prior version of the block in the save volume 488.

6 Consider, for example, a production file system 481 having blocks *a*, *b*, *c*, *d*, *e*, *f*,  
7 *g*, and *h*. Suppose that when the snapshot file system 483 is created, the blocks have  
8 values *a0*, *b0*, *c0*, *d0*, *e0*, *f0*, *g0*, and *h0*. Thereafter, read/write access to the production  
9 file system 481 modifies the contents of blocks *a* and *b*, by writing new values *a1* and *a2*  
10 into them. At this point, the following contents are seen in the clone volume 487 and in  
11 the save volume 488:

12  
13 Clone Volume: *a1*, *b1*, *c0*, *d0*, *e0*, *f0*, *g0*, *h0*

14  
15 Save Volume: *a0*, *b0*

16  
17 From the contents of the clone volume 487 and the save volume 488, it is possible  
18 to construct the contents of the snapshot file system 483. When reading a block from the  
19 snapshot file system 483, the block is read from the save volume 488 if found there, else  
20 it is read from the clone volume 487.

21 FIG. 25 further shows that a snapshot queue 500 maintains respective objects  
22 supporting multiple snapshot file systems 483, 503 created at different respective points  
23 in time from the production file system 481. In particular, the snapshot queue 500

1 includes a queue entry (J+K) at the tail 501 of the queue, and a queue entry (J) at the head  
2 502 of the queue. In this example, the snapshot file system 483, the snapshot volume  
3 484, the delta volume 486, the save volume 488, the bit map 489, and the block map 480  
4 are all located in the queue entry at the tail 501 of the queue. The queue entry at the head  
5 of the queue 502 includes similar objects; namely, the snapshot file system (J) 503, a  
6 snapshot volume 504, a delta volume 505, a save volume 506, a bit map 507, and a block  
7 map 508.

8 The snapshot copy facility 456 may respond to a request for another snapshot of  
9 the production file system 481 by allocating the objects for a new queue entry, and  
10 inserting the new queue entry at the tail of the queue, and linking it to the snapped  
11 volume 485 and the clone volume 487. In this fashion, the save volumes 488, 506 in the  
12 snapshot queue 500 are maintained in a chronological order of the respective points in  
13 time when the snapshot file systems were created. The save volume 506 supporting the  
14 oldest snapshot file system 503 resides at the head 502 of the queue, and the save volume  
15 488 supporting the youngest snapshot file system 483 resides at the tail 501 of the queue.

16 FIG. 26 shows a procedure performed by the snapshot copy facility for writing a  
17 specified block ( $B_i$ ) to the production file system. In step 511, if the snapshot queue is  
18 not empty, execution continues to step 512. In step 512, the bit map at the tail of the  
19 snapshot queue is accessed in order to test the bit for the specified block ( $B_i$ ). Then in  
20 step 513, if the bit is not set, execution branches to step 514. In step 514, the content of  
21 the specified block ( $B_i$ ) is copied from the clone volume to the next free block in the save  
22 volume at the tail of the snapshot queue. Execution continues from step 514 to step 515.  
23 In step 515, the save volume block address ( $S_i$ ) of the free block is inserted into the entry

1 for the block ( $B_i$ ) in the block map at the tail of the queue, and then the bit for the block  
2 ( $B_i$ ) is set in the bit map at the tail of the queue. After step 515, execution continues to  
3 step 516. Execution also continues to step 516 from step 513 if the tested bit is found to  
4 be set. Moreover, execution continues to step 516 from step 511 if the snapshot queue is  
5 empty. In step 516, new data is written to the specified block ( $B_i$ ) in the clone volume,  
6 and then execution returns.

7 FIG. 27 shows a procedure performed by the snapshot copy facility for reading a  
8 specified block ( $B_i$ ) from a specified snapshot file system ( $N$ ). In the first step 521, the  
9 bit map is accessed for the queue entry ( $N$ ) to test the bit for the specified block ( $B_i$ ).  
10 Then in step 522, if the tested bit is set, execution continues to step 523. In step 523, the  
11 block map is accessed to get the save volume block address ( $S_i$ ) for the specified block  
12 ( $B_i$ ). Then in step 524 the data is read from the block address ( $S_i$ ) in the save volume,  
13 and then execution returns.

14 If in step 522 the tested bit is not set, then execution branches to step 525. In step  
15 525, if the specified snapshot ( $N$ ) is not at the tail of the snapshot queue, then execution  
16 continues to step 526 to perform a recursive subroutine call upon the subroutine in FIG.  
17 27 for read-only access to the snapshot ( $N+1$ ). After step 526, execution returns.

18 If in step 525 the snapshot ( $N$ ) is at the tail of the snapshot queue, then execution  
19 branches to step 527. In step 527, the data is read from the specified block ( $B_i$ ) in the  
20 clone volume, and execution returns.

21 FIG. 28 shows the subroutine to resync the primary file system with the secondary  
22 file system, as used in the flowchart of FIG. 21. In a first step 531 in FIG. 28, the delta

1 set identifier (n-x) for the restart point is read from the delta set attribute of the restart  
2 point snapshot, and it is incremented by one to compute (n-x+1). The save volume at the  
3 primary site is searched for the delta set (n-x+1). Execution continues from step 531 to  
4 step 532 if the delta set n-x+1 is found in the save volume at the primary site. In this  
5 case, the primary site should also have all of the delta sets from delta set n-x+1 to delta  
6 set n, and the primary file system can be restored to the state of the restart point in step  
7 532 by an “undo” of the data blocks of these delta sets. The primary file system is  
8 restored with “before images” of these data blocks in order to “undo” the changes. The  
9 delta sets n-x+1 to n contain “after images” of these data blocks. The “before images” of  
10 these data blocks are stored in the secondary site in the snapshot of the secondary file  
11 system at the restart point. Therefore, to restore the primary file system to the state of the  
12 restart point, the snapshot copy facility at the secondary site reads the “before images”  
13 from the “restart point” snapshot, and the secondary site executes a remote copy of these  
14 “before images” to the primary file system.

15 In step 532, the primary site restores the primary file system to the state of the  
16 restarting point by obtaining a list of blocks from the save volume at the primary site,  
17 including the blocks in delta set n-x+1 to delta set n. The primary site sends this list to  
18 the snapshot copy facility at the secondary site. The secondary site retrieves the data of  
19 these blocks from the snapshot at the restart point, and returns the data to the primary site.  
20 The primary site receives these blocks of data and restores them to the primary file  
21 system.

22 In step 533, the snapshot copy facility starts replication to the primary file system,  
23 and creates an intermediate point snapshot (at state delta set n-x+ $\delta$ ), so all new changes

1 made to the secondary file system since the intermediate point are kept and sent to the  
2 primary file system. However, they are not yet played back into the primary file system.  
3 At this point, the primary file system is mounted as a "raw" file system, so it is not  
4 accessible to users.

5 In step 534, the changes made to the secondary file system from the restart point  
6 to the intermediate point are copied from the secondary file system to the primary file  
7 system. These changes are maintained by the snapshot copy facility at the secondary site.  
8 For the snapshot copy facility as shown in FIG. 25, for example, assume that the restart  
9 point is the snapshot J and the intermediate point is the snapshot J+K where K is one. In  
10 this case, the changed blocks are identified in the bit map 507, and the data of the  
11 changed data blocks are found by reading from the intermediate point snapshot file  
12 system. To retrieve these changes, the snapshot copy facility scans the bit map 507 to  
13 find any set bit indicating a change in the save volume, and if the bit is set, the snapshot  
14 copy facility notes the block address ( $B_i$ ). The snapshot copy facility then reads the  
15 snapshot file system (J+K) 483 at this block address ( $B_i$ ). The snapshot copy facility  
16 continues this process until it has finished scanning the bit map 507.

17 In FIG. 28, after step 534, execution continues to step 535. In step 535, the  
18 primary file system is mounted as a read-only UxFS file system, and playback to the  
19 primary file system is started. At this point, the resync is finished, and the primary and  
20 secondary sites are ready for failback.

21 In step 531, if the delta set  $n-x+1$  cannot be found in the save volume from before  
22 failover (459 in FIG. 29) at the primary site (for example, because the save volume was  
23 destroyed during the disaster), then execution branches from step 531 to step 536. In this

1 case, the secondary file system is migrated to the primary file system, and a warning  
2 message is logged and displayed to the system administrator to signal that the full copy  
3 was required. In step 536, the secondary creates a first intermediate point snapshot (at  
4 state delta set  $n-x+\delta$ ), and copies the first intermediate point snapshot to the primary file  
5 system. At this time, read/write access to the secondary file system is permitted on a  
6 priority basis. Then in step 537, the secondary starts replication to the primary file  
7 system, and creates a second intermediate point snapshot, so all new changes made to the  
8 secondary file system since the second intermediate point are kept and sent to the primary  
9 file system by the replication process. However, these new changes are not yet played  
10 back into the primary file system. At this point, the primary file system is mounted as a  
11 "raw" file system, so it is not accessible to users. Moreover, in step 537, the differential  
12 of the first and second intermediate point snapshots (i.e., the changes made to the  
13 secondary file system from the first intermediate point snapshot to the second  
14 intermediate point snapshot) are copied to the primary file system. Then execution  
15 continues to step 535, to mount the primary file system as a read-only UxFS file system,  
16 and to start playback to the primary file system. Also in step 535, the restart snapshot  
17 and the intermediate snapshot can be deleted.

18 The migration method of steps 536 to 537 is used in order to reduce the likelihood  
19 that read/write access to the secondary file system and replication to the primary file  
20 system would be interrupted by the primary and secondary save volumes (274, 275 in  
21 FIG. 5) becoming full. Each save volume functions as a wrap-around buffer. The  
22 primary save volume will store updates received in step 537 from the secondary site until  
23 playback to the primary file system is started in step 535. If the primary save volume

1 would become full in step 537, a TCP flow control signal is sent from the primary site to  
2 the secondary site to suspend transmission of further updates. When transmission of  
3 further updates is suspended, the further updates can be stored in the secondary save  
4 volume (275 in FIG. 5) until the secondary save volume would become full. If the  
5 secondary save volume would become full, then read/write access to the secondary file  
6 system and replication to the primary file system would be suspended until playback to  
7 the primary file system is started in step 535 and updates are removed from the primary  
8 and secondary save volumes.

9 FIG. 29 shows the state of the remote replication system of FIG. 20 during the  
10 resync procedure. At the secondary site, the snapshot copy facility 456 has created the  
11 intermediate point snapshot 457. A list of blocks in the delta sets  $n$  to  $n-x+1$  is obtained  
12 from the save volume 459 at the primary site and sent to the snapshot copy facility 456 at  
13 the secondary site. The snapshot copy facility 456 returns "before images" of the  
14 requested blocks over an IP pipe 458 to the primary file system 451 to restore the primary  
15 file system to the state of the restart point snapshot, and then sends the changes from the  
16 restart point to the intermediate point. The replication service 452 at the secondary site is  
17 replicating changes to the secondary file system since the intermediate point, and is  
18 sending these changes over the IP pipe 460 to the playback service 450. Transmission of  
19 these changes since the intermediate point over the IP pipe 460 is concurrent with  
20 transmission of changes prior to the intermediate point over the IP pipe 458. Once the  
21 primary file system 451 has been synchronized to the intermediate point 457, the  
22 playback service 450 at the primary site is activated to play these changes into the

1 primary file system 451. If the primary site should crash during the copy process of steps  
2 534 or 536, then the copy process can be restarted automatically.

3 FIG. 30 is a flowchart of the subroutine to failback to the primary file system. In  
4 a first step 541, the primary file system is made almost identical to the secondary file  
5 system by the replication process. For example, when the primary file system is within  
6 one delta chunk of the secondary file system, execution continues to step 542. In step  
7 542, the secondary file system is re-mounted as read-only, and the last delta chunk is  
8 copied from the secondary file system to the primary file system in order to synchronize  
9 the primary file system from the secondary file system. Once synchronized, the primary  
10 and secondary are aborted, in order to terminate the processes of replication from the  
11 secondary, playback to the primary, and copying from the secondary to the primary. At  
12 this point, the primary file system contains all the data that the secondary file system had  
13 before the disaster and the additional data that the secondary file system had accumulated  
14 after the disaster. Then in step 543, there is a resumption of the replication of data from  
15 the primary file system and playback to the secondary file system. Then in step 544, the  
16 primary file system is remounted as read/write. Therefore, the recovery process permits  
17 replication to be restarted as it was before the disaster. Some data might have been lost at  
18 failover time since replication is asynchronous. This is dependent on the number "x" of  
19 delta sets that the secondary file system was behind the primary file system when the  
20 disaster occurred at the primary. The less the number "x" of delta sets the closer the final  
21 recovered file system would be to the original file system.

22 A number of commands have been devised for conducting the above recovery  
23 procedure when network file servers (such as shown in FIGS. 1 and 5) are used at the

primary and secondary sites. These commands include a `$fs_replicate-failover` command, a `$fs_replicate -resync` command, a `$fs_replicate-status` command, `$fs_replicate-failback` command, and a `$fs_copy -start` command.

The `$fs_replicate-failover` command has the following format:

```
$fs_replicate -failover <pfs>:cel=<server> <sfs> [-sync] [-now]
```

where `<pfs>` is the primary file system, `<server>` is the server of the primary file system, `<sfs>` is the secondary file system, `[-sync]` is an option that is set when the user wants the failover to happen when both the primary file system and the secondary file system are “in sync” (i.e., identical state), and `[-now]` is an option to specify immediate failover.

The `$fs_replicate-failover` command is executed on the control station for the secondary file system.

When the `[-sync]` option is not set, this command will internally stop the replication and playback of `<sfs>`, making sure that all of the delta sets available on secondary site are re-played. Then it will try to stop the replication on `<pfs>` as well, unless the `[-restartpoint]` option is set. Then `<sfs>` will be remounted “rw”. If possible, `<pfs>` is remounted “ro”. No `[-sync]` option should be specified if the primary control station or server is down. In this case, after failover, some data could be lost.

When the `[now]` option is set, the playback of `<sfs>` is immediately stopped, without replaying any of the delta sets available on the secondary site.

1       When the [-sync] option is specified, the primary control station and the server  
2       must be up, else the command will fail. The primary file system <pfs> is re-mounted as  
3       read-only and a last delta set is created. Once the last delta set is played back to the  
4       secondary, the replication process is stopped on the <sfs>/<pfs>. The secondary file  
5       system <sfs> now contains all the data that the <pfs> had before the failover. The  
6       secondary file system <sfs> is now mounted "rw".

7       FIG. 31 shows a flowchart of execution of the \$fs\_replicate-failover command.  
8       In a first step 551, if the sync option is not set, then execution branches to step 552 to  
9       perform a failover without sync procedure, as described below with reference to FIG. 32.  
10      After step 552, execution is finished. In step 551, if the sync option is set, then execution  
11      continues to step 553. In step 553, if the primary site is not operational, execution returns  
12      reporting a command failure. Otherwise, if the primary site is operational, execution  
13      continues to step 554, to perform a failover with sync procedure, as described below with  
14      reference to FIG. 33.

15      FIG. 32 shows a flowchart of the failover without sync procedure. In step 561, if  
16      possible, the primary file system is remounted as a read-only file system. Then in step  
17      562, the secondary site finishes playback of the delta sets that it had received but had not  
18      played back prior to receipt of the failover command. In step 563, the secondary site  
19      creates a snapshot copy of the secondary file system. This snapshot is called the  
20      restarting point of the file system. In step 564, the secondary file system is remounted as  
21      a read/write file system, and execution of the failover command is finished.

22      FIG. 33 shows a flowchart of the failover with sync procedure. In a first step 571,  
23      the secondary file system is made almost identical to the primary file system by the

1 replication process. For example, once the secondary file system is within one delta set  
2 chunk of the primary file system, execution continues to step 572. In step 572, the  
3 primary file system is remounted as read-only, the last delta chunk is copied from the  
4 primary file system to the secondary file system in order to synchronize the secondary  
5 file system to the primary file system, and then the primary and secondary are aborted to  
6 terminate the processes of replication from the primary, playback to the secondary, and  
7 copying from the primary to the secondary. Then in step 573, the secondary site creates a  
8 snapshot copy of the secondary file system. This snapshot copy is called the restarting  
9 point of the file system. Finally, in step 574, the secondary file system is remounted as  
10 read/write.

11 The `$fs_replicate-resync` command has the following format:

12  
13 `$fs_replicate -resync sfs pfs:cel=server1`  
14

15 where `sfs` is the secondary file system (this is now the source file system), `pfs` is the  
16 primary file system ( this is the file system to be restored, a raw file system mounted on a  
17 server which has its IP interfaces configured), `server1` is the site where a disaster  
18 occurred, and `<ckpt_name>` is the restart point name.  
19

20 The `$fs_replicate-status` command has the following format:  
21

1   \$fs\_replicate -status <pfs> <sfs>:cel=<remote\_server>

2

3   where <pfs> is the primary file system, <sfs> is the secondary file system, and  
4   <remote\_server> is the server of the secondary site. This command can be used on the  
5   control station for the primary site, and also on the control station for the secondary site.  
6   Typically this command is used on the primary site when the user wants to check the  
7   status of the replication process, to determine when it is reasonable to failback. It is  
8   considered reasonable to failback when the difference between the primary file system  
9   and the secondary file system is very small. To assist the user, the primary site responds  
10   to this command by providing an indication of the number of blocks that are different  
11   between the primary file system and the secondary file system.

12       Following is an example of the information provided when the \$fs\_replicate-  
13   status command is executed:

```
14       [nasadmin@cel_sfs nasadmin]$ fs_replicate -info sfs10
15       -v 128
16       id                               = 59
17       name                             = sfs10
18       type                             = replication
19       current_delta_set               = 16
20       current_block_number            = 1088
21       current_transfer_rate           = 17.0667 blocks/second
22       avg_transfer_rate               = 25.7444 blocks/second
23       id                               = 31
24       name                             = pfs10:cel_ip2
```

1 type = playback

2 current\_delta\_set = 16

3

4 outstanding delta sets:

5 Delta Source\_create\_time Blocks

6 -----

7 15 11/26 20:11:56 1

8 | Source | Destination

9 Delta|Create Dur Blocks|Playback Dur Blocks Dsin

10 | Time | Time Group

11 -----

12 15 11/26 20:11:56 0 1

13 14 11/26 20:01:56 0 1

14 13 11/26 19:51:56 0 1 11/26 21:08:48 0 1 2

15 12 11/26 19:41:56 0 1

16 11 11/26 19:31:56 0 1 11/26 20:48:48 0 1 2

17 10 11/26 19:21:56 0 1

18 9 11/26 19:11:56 0 1 1 /26 20:28:28 0 1 2

19 8 11/26 18:52:40 0 1 11/26 20:08:09 0 1 1

20 7 11/26 18:42:40 0 1

21 6 11/26 18:32:40 0 1 11/26 19:48:09 0 1 2

22 5 11/26 18:22:39 0 1

23 4 11/26 18:12:39 0 1 11/26 19:28:08 0 1 2

24 3 11/26 18:02:39 0 4 11/26 19:07:23 0 4 1

1  
2 The \$fs\_replicate-failback command has the following format:

3  
4 `$fs_replicate -failback <pfs> <sfs>:cel=<remote_server>`

5  
6 where <pfs> is the primary file system (a UxFS file system mounted as “ro”), and <sfs>  
7 is a secondary file system. This command is used when the user wants to failback to the  
8 primary site (after re-build phase). No data is lost and the replication will continue but it  
9 will be reversed (the primary file system, which runs playback, will then run replication  
10 and vice-versa for the secondary file system). This command can also be used to change  
11 the direction of replication, in the absence of any failure or disruption.

12 The \$fs\_replicate-failback command is executed on the site that the replication  
13 service is running (i.e., the site having the file system that is mounted as read/write).  
14 Typically, the \$fs\_replicate-failback command is executed on the primary site after it has  
15 been rebuilt. Then the primary file system is a raw file system that is restored from the  
16 secondary file system. The secondary file system is re-mounted read write. Replication  
17 is stopped on the secondary and playback is stopped on the primary. After the primary  
18 file system and the secondary file systems are in sync, the primary file system is  
19 converted to a UxFS file system and re-mounted as read/write. Replication is turned  
20 back on and the flow of data is now from the primary file system to the secondary file  
21 system. Playback is started on the secondary file system. This command can be used to  
22 “swap” primary and secondary without having to do a full copy and without having to

1 restart replication. On error, it should be possible to re-issue the `fs_replicate -failback`  
2 `command` and proceed if the error condition has been cleared.

3 The `$fs_copy-start` command has the following format:

4  
5 `$fs_copy -start <fs_ckpt2> <sfs>:cel=<remote_server> -fromfs <fs_ckpt1>`

6  
7 where `<fs_ckpt1>` is the primary file system ckpt1, `<fs_ckpt2>` is the primary file system  
8 ckpt2, `<sfs>` is the secondary file system, and `<remote_server>` is the server of the  
9 secondary site. This command is used internally by the `fs_replicate-resync` command. It  
10 could also be used for other purposes, for example, as described below with reference to  
11 FIGS. 36-38. This command copies the delta from `fs_ckpt1` to `fs_ckpt2` over to the  
12 secondary file system, knowing that the state of the secondary file system was equal to  
13 `fs_ckpt1`.

14 The following is an example of how the commands can be used to recover from a  
15 disaster.

- 16 1. After the disaster, decide the option to use for failover.
- 17 2. `$fs_replicate -failover pfs:cel=cel_pfs sfs`
- 18 3. Verify that sfs can accept read/write operation.
- 19 4. Initialize replication from sfs to pfs:

20 On PFS control station:

21 `server_mount server_2 -o ro pfs /pfs`

22 On SFS control station:

23 `$fs_replicate -resync pfs:cel=cel_pfs sfs`

- 1           5.       Verify that replication from sfs to pfs is running without a problem.
- 2           6.               \$fs\_replicate -failback sfs pfs:cel=cel\_pfs
- 3           7.       Verify that replication from pfs to sfs is running without a problem.

#### 4

#### 5       Replication of Snapshots Using IP File System Copy Differential

6           As described above with reference to steps 534 and 537 of FIG. 28, a file system  
7       copy differential was used to replicate an intermediate point snapshot during recovery of  
8       a file system from a disaster. A general-purpose file system copy differential command  
9       (\$fs\_copy-start) was also introduced.

10          FIG. 34 shows a flowchart of a procedure for the \$fs\_copy -start command. Here  
11       L is the index of the snapshot fs\_ckpt1, and M is the index of the snapshot fs\_ckpt2. In a  
12       first step 581, M is compared to L, and execution branches to step 582 to return an error  
13       if M is not greater than L. Otherwise, execution continues to step 583:

14          The following steps execute the \$fs\_copy -start command by scanning one or  
15       more of the bit maps of the snapshot copy facility of FIG. 25. Each bit map indicates  
16       whether or not each block of a respective snapshot changed from the snapshot to the next  
17       snapshot. Therefore, by examining the bit maps for the snapshots L to M-1, it is  
18       determined whether or not each block has changed from snapshot L to M. If so, the new  
19       value of the block is read from the snapshot M. The bit maps are scanned in block order  
20       so that the new blocks are read from the snapshot M of the primary (i.e., the source) file  
21       system and transmitted to the secondary (i.e., the destination) file system in the order of  
22       their block indices ( $B_i$ ).

1 In step 583, the block index ( $B_i$ ) is set to zero. In step 584, a snapshot index ( $I$ ) is  
2 set to  $L$ . In step 585, the bit map for the snapshot ( $I$ ) is indexed with the block index ( $B_i$ )  
3 to determine whether or not the block was changed between snapshot ( $I$ ) and snapshot  
4 ( $I+1$ ). If in step 585 the bit for the block ( $B_i$ ) is not set in the bit map for the snapshot ( $I$ ),  
5 then no such change occurred, and execution continues to step 586. In step 586, if the  
6 snapshot index  $I$  is not equal to  $M-1$ , then execution continues step 587. In step 587, the  
7 snapshot index  $I$  is incremented by one. After step 587, execution loops back to step 585.

8 In step 585, if the bit for the block ( $B_i$ ) is set in the bit map for the snapshot ( $I$ ),  
9 then execution branches to step 589. In step 589, the snapshot ( $M$ ) is read to get the new  
10 data for the block ( $B_i$ ). Execution continues to step 590 of FIG. 35. In step 590, the  
11 block index ( $B_i$ ) and the new data for the block ( $B_i$ ) from the snapshot ( $M$ ) are returned  
12 for copying into the secondary file system. Then in step 591, if the block index ( $B_i$ ) is at  
13 the end of the production volume, then execution is finished. If not, execution loops back  
14 to step 588 of FIG. 34. In step 588, the block index ( $B_i$ ) is incremented by one, and  
15 execution loops back to step 584, to continue scanning for the next value of the block  
16 index.

17 In step 586 of FIG. 34, if  $I$  is equal to  $M-1$ , then execution continues to step 591  
18 of FIG. 35.

19 It should be understood that the flowchart of FIGS. 34-35 represents a program  
20 executed by at least one processor in a data storage system such as a network file server.  
21 The processor, for example, is a data mover computer (e.g., 232 in FIG. 1). The program,

1 for example, is initially contained in a program storage device such as a floppy disk (e.g.,  
2 238 in FIG. 1) and down-loaded into storage of the data mover computer.

3 The program in the flowchart of FIGS. 34-35 has an inner loop including steps  
4 585, 586, 587 that indexes the snapshots L to snapshot M-1. This sequence includes the  
5 snapshot L and the snapshots that are both younger than the snapshot L and older than the  
6 snapshot M. The program in the flowchart of FIGS. 34-35 has an outer loop including  
7 steps 584, 585, 586, 591, and 588 that indexes the blocks. When a bit in the indexed bit  
8 map is found to be set in step 585, the inner loop is exited to return the block index ( $B_i$ )  
9 and the data in the snapshot M for block ( $B_i$ ).

10 The snapshot copy differential has been described above for facilitating recovery  
11 of a file system after a disaster. The snapshot copy differential can also be used for wide-  
12 area distribution of updates on an as-needed basis. This reduces network traffic for the  
13 case where a client has an old local version of a file system and needs a new version of  
14 the file system. A new local version of the file system can be constructed by copying the  
15 appropriate changes into the old local version of the file system.

16 FIG. 36, for example, shows a block diagram of a data network in which snapshot  
17 deltas are transmitted over a wide-area network 626 from a network file server 627 to a  
18 local file server 624 in order to update a local file system 625 as needed. The local file  
19 server 624 services local clients 621, 622, and 623. When a client needs a more recent  
20 version of the file system, and the local file system 625 is not the most recent version,  
21 then the local file server may request a specified version from the network file server 627,  
22 or the client may request the most recent version available.

1           The network file server 627 has a snapshot copy facility 628 storing multiple  
2 snapshots 629, 630. If the local file system 625 in the local file server 624 is one of the  
3 multiple snapshots, then the network file server 627 may respond to a request from the  
4 local file server 624 by obtaining from the snapshot copy facility 628 a snapshot copy  
5 differential that would contain all of the updates necessary to convert the local file system  
6 624 to a more recent snapshot of the file system. In the usual case, the local file server  
7 624 would request all of the updates necessary to convert the local file system 625 to the  
8 most recent snapshot copy. However, it is also possible for the local file server 624 to  
9 request the updates for a specified version that would not necessarily be the most recent  
10 snapshot copy.

11           FIGS. 37 and 38 show a flowchart of a procedure for the replication of the most  
12 recent snapshot in the system of FIG. 36 using the snapshot copy differential. In a first  
13 step 641, the client requests access to the file in the local file system in the local file  
14 server. In step 642, the local file server accesses attributes of the local file system and  
15 finds that it is version (Q) of a local updatable file system, last updated at time (Tu) from  
16 a network file server having a network identifier (NETID). The local file server uses the  
17 time of last update (Tu) to determine that it is time to check for a more recent version,  
18 and sends an update request to the network file server (NETID). For example, the file  
19 system has an attribute specifying a minimum update interval (Tmin), and it is not time to  
20 request an update unless the minimum update interval has elapsed since the time of the  
21 last update. The request specifies the version (Q) already in the local file server. In  
22 response, the network file server accesses the snapshot copy facility to find the oldest  
23 version (J) and the youngest version (J+K) stored in the network file server. In step 645,

1 Q is compared to J+K. If  $Q=J+K$ , then execution branches to step 646 of FIG. 38. In  
2 step 646, the network file server returns a message that no more recent version is  
3 available. The local file server resets the time of last update ( $T_u$ ) to the current time, and  
4 accesses the local version (Q) for the client. Upon completion of step 646, the snapshot  
5 replication process is finished.

6 If in step 645 Q is not equal to J+K, then execution continues to step 647 of FIG.  
7 38. In step 647, Q is compared to K. If Q is less than K, then execution branches to step  
8 648. In this case, the version (Q) is not a snapshot in the snapshot copy facility because  
9 the version (Q) is too old. In step 648, the network file server copies the youngest  
10 snapshot version (J+K) to the local file server. The local file server replaces the local  
11 version (Q) with the new version (J+K), resets the time of last update ( $T_u$ ) to the current  
12 time, and accesses the new local version (J+K) for the client. Upon completion of step  
13 648, the snapshot replication process is finished.

14 In step 647, if Q is not less than K, then execution continues to step 649. In step  
15 649, the network file server does a file system copy snapshot delta  $\langle Q \rangle$  to  $\langle J+K \rangle$  of  
16 blocks into the local version (Q) to convert it into the youngest snapshot version (J+K).  
17 The local file server resets the time of last update ( $T_u$ ) to the current time, and accesses  
18 this local snapshot version (J+K) for the client. Upon completion of step 649, the  
19 snapshot replication process is finished.

20 In a preferred snapshot copy facility, as described below with reference to FIGS.  
21 41 to 46, there is kept a meta bit map for each snapshot copy for indicating blocks of the  
22 production file system that are not used in the snapshot copy. Further details of such a  
23 snapshot copy facility are described in Philippe Armangau, et al., "Data Storage System

1 Having Meta Bit Maps for Indicating Whether Data Blocks are Invalid in Snapshot  
2 Copies,” U.S. Patent Application Ser. 10/213,241 filed Aug. 6, 2002, incorporated herein  
3 by reference. The snapshot copy facility maintains the meta bit maps in order to store the  
4 “before image” of a block in the save volume at the tail of the snapshot queue only when  
5 the block is being written to and the “before image” is needed for responding to any  
6 request for reading a snapshot copy. This reduces the number of blocks that are stored in  
7 the save volumes. However, in this case, the bit map for each snapshot (L) indicates  
8 whether or not a block has been stored in the save volume for the snapshot (L), and no  
9 longer will indicate all of the blocks that have been changed after snapshot (L) and before  
10 snapshot (L+1). In particular, if a block was not in use for snapshot (L), and was written  
11 to after snapshot (L) and before snapshot (L+1), then the “before image” of the block will  
12 not be written into the save volume for snapshot (L).

13 When it is known that a block is not used in the snapshot copy (M), then there is  
14 no need for the snapshot copy facility to return the block when responding to a request  
15 for the snapshot delta of snapshot <L> to snapshot <M>. Therefore, for the preferred  
16 snapshot copy facility, it is desirable to modify the procedure of FIG. 34 in order use the  
17 information in the meta bit map for the snapshot <M>. In this case, the procedure of  
18 FIG. 34 should also be modified to account for the fact that the save volumes no longer  
19 store the “before images” for all of the blocks that may have changed between the  
20 successive snapshot copies.

21 FIG. 39 shows how the flowchart of FIG. 34 can be modified for use with the  
22 preferred snapshot copy facility of FIGS. 41 to 46. Steps 651, 652, and 653 of FIG. 39

1 are similar to steps 581, 582, and 583 of FIG. 34. Step 654 of FIG. 39 is similar to step  
2 588 of FIG. 34, and step 658 of FIG. 39 is similar to step 589 of FIG. 34.

3 In FIG. 39, after step 653 or step 654, execution continues to step 655. In step  
4 655, if the block ( $B_i$ ) is not in use in the snapshot ( $M$ ), then execution branches to step  
5 591 of FIG. 35. Therefore, a block not in use in the snapshot ( $M$ ) is not returned in  
6 response to the command to copy the snapshot delta  $\langle L \rangle$  to  $\langle M \rangle$ . If the block ( $B_i$ ) is in  
7 use in the snapshot ( $M$ ), then execution continues from step 655 to step 656. In step 656,  
8 if the block ( $B_i$ ) is in any of the save volumes ( $L$ ) to ( $M-1$ ), then execution continues to  
9 step 658 to read the block ( $B_i$ ) from the snapshot ( $M$ ) in order to return this version of the  
10 block ( $B_i$ ) in response to the command to copy the snapshot delta  $\langle L \rangle$  to  $\langle M \rangle$ . In step  
11 656, if the block ( $B_i$ ) is in any of the save volumes ( $L$ ) to ( $M-1$ ), then execution branches  
12 to step 657. In step 657, if the block ( $B_i$ ) is in use in all of the snapshots ( $L$ ) to ( $M-1$ ),  
13 then execution branches to step 591 of FIG. 35. In this case, the block ( $B_i$ ) did not  
14 change from snapshot ( $L$ ) to snapshot ( $M$ ), because if it did, a "before image" would have  
15 been stored in one of the save volumes ( $L$ ) to ( $M-1$ ). In step 657 of FIG. 39, if it is not  
16 true that the block ( $B_i$ ) is in use in all of the snapshots (i.e., it is true that the block ( $B_i$ ) is  
17 not in use in at least one of the snapshots ( $L$ ) to ( $M-1$ )), then execution continues from  
18 step 657 to step 658 to read the block ( $B_i$ ) from the snapshot ( $M$ ) in order to return this  
19 version of the block ( $B_i$ ) in response to the command to copy the snapshot delta  $\langle L \rangle$  to  
20  $\langle M \rangle$ . In this case, it is possible that the block changed from snapshot ( $L$ ) to snapshot  
21 ( $M$ ) despite the fact that the block ( $B_i$ ) is not in any of the save volumes ( $L$ ) to ( $M-1$ ).

FIG. 40 shows a preferred implementation of the procedure of FIG. 39. Steps 661 to 665 of FIG. 40 are similar to steps 651 to step 655 of FIG. 39. In step 665, the meta bit map for snapshot (M) has a value for the block (B<sub>i</sub>) indicating whether or not the block (B<sub>i</sub>) is in use for the snapshot (M). In particular, a value of 1 indicates that the block (B<sub>i</sub>) is in use for the snapshot (M). Steps 666, 669, and 670 of FIG. 40 are similar to step 584, 586 and 587 of FIG. 34. From step 666 or step 670, execution continues to step 667.

In step 667, if the bit map for snapshot (I) has a value of 1 for the block (B<sub>i</sub>), then execution continues to step 671 to read the snapshot (M) to get data for the block (B<sub>i</sub>), in order to return the data in response to the command to copy the snapshot delta <L> to <M>. In this case, the save volume for block (I) includes a "before image" for the block (B<sub>i</sub>). Otherwise, if the bit map for snapshot (I) does not have a value of 1 for the block (B<sub>i</sub>), execution branches from step 667 to step 668. In step 668, if the meta bit map for the snapshot (I) does not have a value of 1 for the block (B<sub>i</sub>), execution continues to step 671 to read the snapshot (M) to get data for the block (B<sub>i</sub>), in order to return the data in response to the command to copy the snapshot delta <L> to <M>. In this case, the block (B<sub>i</sub>) is not in use in the snapshot (I). Otherwise, if the meta bit map for the snapshot (I) has a value of 1 for the block (B<sub>i</sub>), execution continues to step 669.

#### **Maintenance of Meta Bit Maps in the Snapshot Copy Facility**

In the above description of the snapshot copy process, and in particular with respect to FIG. 25, it was assumed that the original contents of a block of the production file system must be saved to the most recent save volume before the contents of the block

1 are modified by a write access to the production file system. In practice, however, the  
2 original contents are often invalid, and therefore need not be saved. For example, many  
3 applications start with an empty file, and the file increases in size as data is written to the  
4 file. In some of these applications, the file rarely decreases in size. However, storage for  
5 the file may be released when the file is deleted from the file server, for example, when  
6 the file is transferred to archival storage. In some applications, the extent of a file may be  
7 dynamically decreased concurrent with read/write access to the file.

8       There are significant advantages to identifying when read/write access to the  
9 production file system is about to modify the contents of an invalid data block. If this can  
10 be done in an efficient manner, then there can be a decrease in the access time for write  
11 access to the production file system. A write operation to an invalid block can be  
12 executed immediately, without the delay of saving the original contents of the data block  
13 to the most recent save volume at the tail of the snapshot queue. Moreover, there is a  
14 saving of storage because less storage is used for the save volumes. There is also a  
15 decrease in memory requirements and an increase in performance for the operations upon  
16 the snapshot file systems, because smaller bit and block hash indices can be used, and the  
17 reduced amount of storage for the snapshots can be more rapidly restored to the  
18 production file system, or deallocated for re-use when snapshots are deleted.

19       An efficient way of identifying when read/write access to the production file  
20 system is about to modify the contents of an invalid data block is to use a meta bit map  
21 having a bit for indicating whether or not each allocated block of storage in the  
22 production file system is valid or not. For example, whenever storage is allocated to the  
23 production file system by the initial allocation or the extension of a clone volume, a

1 corresponding meta bit map is allocated or extended, and the bits in the meta bit map  
2 corresponding to the newly allocated storage are initially reset.

3 FIG. 41 shows a procedure for writing a specified block (Bi) to the production file  
4 system when there is a meta bit map for indicating invalid data blocks in the production  
5 file system. In a first step 681, a queue pointer is set to point to the queue entry at the tail  
6 of the snapshot queue. Next, in step 682, the bit map in this snapshot queue entry is  
7 accessed to test the bit for the specified block (Bi). Next, in step 683, if the tested bit is  
8 found to be set, then execution continues to step 684. In step 684, new data is written to  
9 the specified block (Bi) in the clone volume, and then execution returns.

10 In step 683, if the bit in the bit map is not set, then execution branches to step 685.  
11 In step 685, the meta bit map in the snapshot queue entry is accessed to test the bit for the  
12 specified block (Bi). Then, in step 686, execution continues to step 687 if this bit is  
13 found to be set. In step 687, the content of the block (Bi) is copied from the clone  
14 volume to the next free block in the save volume at the tail of the snapshot queue. In step  
15 688, an entry for the block (Bi) is inserted into the block map at the tail of the snapshot  
16 queue, and then the bit for the block (Bi) is set in the bit map at the tail of the snapshot  
17 queue. Execution continues from step 688 to step 684, to write new data to the block (Bi)  
18 in the clone volume.

19 In step 686, if the tested bit is found not to be set, then execution branches to step  
20 689. If the queue pointer is pointing to the head of the queue, then execution branches to  
21 step 684, to write new data to the block (Bi) in the clone volume. Otherwise, if the queue  
22 pointer is not pointing to the head of the snapshot queue, then execution continues to step

1 690 to advance the queue pointer to the next snapshot queue entry toward the head of the  
2 snapshot queue. After step 690, execution loops back to step 682.

3 FIG. 42 shows an organization of the snapshots in the network file server when a  
4 respective meta bit map 692, 693 is maintained for each snapshot in addition to the meta  
5 bit map 691 for the production volume. It is desired to maintain a respective meta bit  
6 map for each snapshot so that whenever the production file system is restored with a  
7 snapshot file system, the meta bit map for the production file system can be restored with  
8 the meta bit map for each snapshot. For example, a meta bit map of the production file  
9 system is included in the production file system, so that whenever a snapshot copy of the  
10 production file system is created, a snapshot copy of the meta bit map is also created.  
11 Consequently, when the production file system is restored with a snapshot, the meta bit  
12 map of the production volume is replaced with the meta bit map of the snapshot.

13 As shown in FIG. 42, a meta bit map 691 is linked to the production file system  
14 481 for indicating invalid blocks in the production file system. Each entry in the  
15 snapshot queue 500 includes a respective meta bit map linked to the snapshot file system  
16 in the entry. For example, the queue entry (J+K) at the tail 501 of the queue has a meta  
17 bit map 692 linked to the snapshot file system 483, and the queue entry (J) at the head  
18 502 of the queue includes a meta bit map 693 linked to the delta volume 505.

19 To reduce the memory and storage requirements for maintaining the bit maps 694,  
20 696 and block maps 695, 697, the each bit map is organized as a set of pages indexed by  
21 a page table, and the each block map is organized as a set of hash lists indexed by a hash  
22 table. The bit maps and block maps 694, 695 at the queue entry (J+K) at the tail of the  
23 queue are initially created in a random access memory cache and written back to storage

1 of the save volume 488 when a next snapshot of the production file system 481 is created.  
2 Thereafter the bit maps and block maps can be deallocated from the random access  
3 memory in order to free up cache memory, and later staged from storage to the cache  
4 memory when needed for read access to their respective snapshots.

5 FIG. 43 shows that the bit map 694 is organized as a page table 701 indexed by a  
6 set of most significant bits (MSB) of the block index ( $B_i$ ), and pages 702 and 703 of the  
7 bit map linked to respective entries of the page table. The page table 701 includes a set  
8 of entries, each of which is either zero, indicating that the entire page is zero and  
9 therefore the page is not stored in storage or random access memory, or is a pointer to a  
10 page of the bit map.

11 FIG. 44 shows that the block map 695 is organized as a hash table 708 indexed by  
12 a hashing of the block index ( $B_i$ ), and a set of hash lists 709 linked to respective entries  
13 of the hash table 708. Each non-zero entry in the hash table 708 points to a respective  
14 one of the hash lists 709. Each entry in each hash list includes a block address ( $B_i$ ) to a  
15 block in the clone volume, a corresponding block address ( $S_i$ ) of the block in the save  
16 volume, and a value that is either zero indicating the end of the has list, or a pointer to the  
17 next entry in the list.

18 It is not necessary to retain the bit map 694 for the snapshot because the block  
19 map 695 can be accessed to provide the same information that is contained in the bit map  
20 694. In particular, the block map 695 can be accessed simply to determine whether or not  
21 a specified block ( $B_i$ ) is found in the block map, and therefore can be found in the save  
22 volume for the corresponding snapshot. However, the bit map 694 can be accessed more  
23 quickly to provide this information. Therefore, for applications such as backup where the

1 snapshots would be accessed very infrequently, it may be desirable to conserve storage  
2 by discarding the bit map for the snapshot at the tail of the queue once the next snapshot  
3 is created. Otherwise, for an application such as described above with respect to FIGS.  
4 36-38 where the snapshots would be accessed frequently, the bit map for the snapshot at  
5 the tail of the snapshot queue may be written to storage and maintained in a random  
6 access cache memory as needed.

7 FIG. 45 shows that the meta bit map 691 has a respective bit corresponding to  
8 each block in the clone volume, and in this example, each bit in the meta bit map  
9 corresponds to one and only one block in the clone volume. The meta bit map 691  
10 includes a series of words, each with a multiple of M bits. In this example, a bit having a  
11 value of zero indicates a corresponding block that is invalid, and a bit having a value of  
12 one indicates a corresponding block that is valid.

13 The meta bit map, however, may have a granularity greater than one block per bit.  
14 For example, each bit in the meta bit map could indicate a range of block addresses,  
15 which may include at least some valid data. The benefit to the increased granularity is a  
16 reduced size of the meta bit map at the expense of sometimes saving invalid data to the  
17 save volume. For example, FIG. 46 shows the interpretation of a meta bit map 691'  
18 having a granularity of two blocks per bit. Each bit is set if any one of the two  
19 corresponding blocks is valid, or conversely, each bit is clear only if neither of the two  
20 corresponding blocks is valid. In this case, the block address can be converted to a bit  
21 address by an integer division by two, for example, by an arithmetic right shift of the  
22 block address by one bit position.

1           In view of the above, there has been described an asynchronous remote copy  
2 system including a primary file system at a primary site and a secondary file system at a  
3 secondary site. The primary site keeps a list of data blocks for which changes have been  
4 made to the primary file system and transmitted to the secondary file system. When the  
5 primary site becomes inoperative, the secondary site begins read/write access to the  
6 secondary file system and keeps a snapshot copy of the restart point. Once the primary  
7 site becomes operative, the primary file system is restored to the state of the secondary  
8 file system at the restart point by using the list of data blocks for obtaining from the  
9 snapshot copy the data of the blocks for which changes had been made to the primary file  
10 system after the restart point. Then the primary file system is synchronized to the  
11 secondary file system, and read/write access is switched back from the secondary file  
12 system to the primary file system. The primary file system can be synchronized by  
13 making a snapshot copy of the secondary file system at an intermediate point, beginning  
14 a process of asynchronous remote copy of changes made to the secondary file system  
15 since the intermediate point to the primary file system, and transmitting to the secondary  
16 file system the changes made to the secondary file system from the restart point to the  
17 intermediate point. Once the changes made to the secondary file system from the restart  
18 point to the intermediate point have been written into the primary file system, the changes  
19 made to the secondary file system since the intermediate point are played back into the  
20 primary file system.

21           The asynchronous remote copy system can replicate read-only copies using a  
22 standard network protocol such as the Internet Protocol. Therefore, the relative cost of  
23 backup is reduced because a dedicated link is not needed between the primary site and

1 the secondary site, and the secondary site can also be used for local access to the read-  
2 only copies.

3